A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device

ZVIKA BRAKERSKI, Weizmann Institute of Science, Israel

PAUL CHRISTIANO, OpenAI, USA

URMILA MAHADEV, California Institute of Technology, USA

UMESH VAZIRANI, UC Berkeley, USA

THOMAS VIDICK, California Institute of Technology, USA

We consider a new model for the testing of untrusted quantum devices, consisting of a single polynomial time bounded quantum device interacting with a classical polynomial time verifier. In this model, we propose solutions to two tasks—a protocol for efficient classical verification that the untrusted device is “truly quantum” and a protocol for producing certifiable randomness from a single untrusted quantum device. Our solution relies on the existence of a new cryptographic primitive for constraining the power of an untrusted quantum device: post-quantum secure trapdoor claw-free functions that must satisfy an adaptive hardcore bit property. We show how to construct this primitive based on the hardness of the learning with errors (LWE) problem.

CCS Concepts: • **Theory of computation** → **Interactive proof systems**;

Additional Key Words and Phrases: Test of quantumness, certified randomness, quantum interactive proofs

ACM Reference format:

Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. 2021. A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device. *J. ACM* 68, 5, Article 31 (August 2021), 47 pages.

<https://doi.org/10.1145/3441309>

1 INTRODUCTION

The testing of quantum devices, besides being a pressing practical challenge, touches on foundational questions in quantum computational complexity. The classical verifier of such a device is necessarily at a disadvantage due to the exponential power of quantum systems, and the laws of quantum mechanics severely limit the amount of information that can be accessed in principle. Nevertheless, a sequence of results has shown that it is possible to verify the correctness of untrusted quantum devices (also referred to as provers) in a variety of settings, including certifiable

Authors' addresses: Z. Brakerski, Weizmann Institute of Science, Rehovot 7610001, Israel; email: zvika.brakerski@weizmann.ac.il; P. Christiano, OpenAI, USA; U. Mahadev and T. Vidick, California Institute of Technology, Pasadena CA 91106, USA; emails: mahadev@berkeley.edu, vidick@caltech.edu; U. Vazirani, UC Berkeley, Berkeley CA 94701, USA; email: vazirani@cs.berkeley.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

0004-5411/2021/08-ART31 \$15.00

<https://doi.org/10.1145/3441309>

random number generation, quantum key distribution, and quantum computation. These results have been established in two models: In the first, the classical verifier is augmented with the ability to prepare a sequence of quantum states on small numbers of qubits and transmits them to the quantum device [3, 4, 12, 18], and in the second, the classical verifier interacts with two non-communicating quantum devices that share entanglement [16, 38, 44].

In this article, we consider a new model, in which a purely classical verifier interacts with a single, polynomial time bounded quantum machine. The restriction to an efficient quantum device allows the verifier to leverage post-quantum cryptography, i.e., cryptographic primitives that can be implemented efficiently on a classical computer but that cannot be broken by any efficient quantum computer.

In this model, we propose solutions to two basic tasks: how to efficiently verify that an untrusted device is “truly” quantum, and how to generate certifiably random strings from a single untrusted quantum device. The first task is also referred to as “quantum supremacy,” and existing protocols for this [1, 2, 7, 9, 10] rely on exponential time classical verification using a classical supercomputer. By contrast, our qubit certification test below provides a proof of quantumness that can be verified by a classical verifier in polynomial time. There has also been considerable research into certifiable random number expansion from quantum devices [6, 16, 31, 35, 43], including experimental demonstrations [8, 35]. However, all prior works have focused on the setting where there are multiple quantum devices that share entanglement and where the randomness certification relies on the violation of a Bell inequality.

The core of the difficulty in interacting with untrusted quantum devices lies in enforcing a qubit structure in the device’s operations, i.e., that the quantum device actually holds qubits and is performing measurements on them to respond to the verifier’s queries. In the two previous models of testing quantum devices, this issue was handled in two different ways. For slightly quantum verifiers, the verifier could simply send qubits to the prover, encoded in such a way that the prover was forced to work with only those qubits. In the model of two entangled quantum devices, Bell inequality violations were used to prove that the two devices must share Bell states and measure them in the X and Z basis as requested. Our work relies on post-quantum cryptography to enforce qubits. Roughly, our qubit certification protocol enables the quantum device (the prover) to create a qubit in the state $\frac{1}{\sqrt{2}}|0\rangle + (-1)^b \frac{1}{\sqrt{2}}|1\rangle$, where the bit b is computationally hidden from the prover. With knowledge of the trapdoor for the post-quantum cryptosystem, the classical verifier can compute b and use it to verify that the prover actually holds the above state, thereby gaining leverage over the quantum prover.

Our certifiable randomness protocol uses the qubit certification protocol as a subroutine and provides an information-theoretic guarantee about the random string output by the untrusted quantum device. The guarantee is stronger than computational pseudorandomness, which is easily achievable under standard cryptographic assumptions, since the verifier starts with a short uniformly random seed. It is illuminating to understand how an information-theoretic guarantee could even be connected to the computational assumptions about the device. We imagine that there is an adversary with unbounded computing power and an unboundedly large quantum register E , which may be entangled with the quantum device register D . The guarantee can now be expressed as saying that the unbounded adversary, who is allowed to design the quantum device and to perform an arbitrary measurement on the register E , cannot distinguish the output of the protocol from a uniform sequence of bits, provided the device is unable to break the post-quantum cryptography during the execution of the protocol.

A qubit certification test. The specific cryptographic primitive we rely on is a post-quantum secure trapdoor claw-free (in short, TCF) family of function pairs $f_0, f_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$, the

-
- (1) The verifier generates a TCF pair, along with a trapdoor, and sends just the function pair to the prover.
 - (2) The prover returns an image y of the TCF pair.
 - (3) The verifier challenges the prover by randomly asking for either a preimage of y , or a bit c and an n -bit string d such that $d \cdot (x_0 \oplus x_1) = c$.
 - (4) The prover measures in standard or Hadamard basis to return the requested output and the verifier checks the validity by using the trapdoor to compute the two preimages x_0, x_1 of y .
-

Fig. 1. The quantum certification protocol.

post-quantum analogue of a notion introduced by Goldwasser, Micali, and Rivest in the context of digital signatures [22]. A TCF pair is a pair of functions that are injective, with the same image, and satisfy the following property. With knowledge of a secret trapdoor it is possible to efficiently (classically) compute the two preimages x_0 and x_1 of a given y ($f_0(x_0) = f_1(x_1) = y$), but without the trapdoor, there is no efficient quantum algorithm that can compute such a triple (x_0, x_1, y) , referred to as a *claw*, for any y .

While the quantum device cannot compute a claw, nevertheless it can simultaneously hold an image y as well as a superposition

$$\frac{1}{\sqrt{2}}(|0\rangle|x_0\rangle + |1\rangle|x_1\rangle) \quad (1)$$

over the two preimages of y , simply by evaluating f on a uniform superposition over all inputs and measuring the image y . If the quantum device were to measure the above state in the standard basis, then it would obtain a random preimage, x_0 or x_1 . This is not particularly interesting, since a classical machine could sample from the same distribution by first sampling a random bit b and string x and then computing $y = f_b(x)$. To take advantage of the fact that the preimages x_0, x_1 are stored in superposition, the quantum device can instead perform a Fourier (Hadamard) basis measurement on all but the first qubit of the state, yielding a string $d \in \{0, 1\}^n$. At this point, we are back to the state mentioned earlier; the quantum device currently holds, for $c = d \cdot (x_0 \oplus x_1)$,

$$\frac{1}{\sqrt{2}}(|0\rangle + (-1)^c|1\rangle). \quad (2)$$

A Fourier measurement of the single qubit state above will yield the bit c .

This Fourier measurement is the aspect that separates quantum and classical devices; intuitively, the output (d, c) should be hard to reproduce in the classical setting, as it is dependent on both elements x_0 and x_1 in the superposition and the claw-free property implies that it is computationally intractable to hold both x_0 and x_1 simultaneously. This suggests the qubit certification test, between a classical verifier and a quantum prover, written in Figure 1.

The quantum prover can successfully answer either challenge in the qubit certification protocol by measuring the state in Equation (1) in the standard or Hadamard basis. By contrast, we would like to argue that no classical algorithm can succeed at this task. This is counter-intuitive, as ultimately our proof must rely on the security of the TCF, which applies equally to classical and quantum attacks. The crux of the proof is that classical computations can be rewound, while quantum measurements cannot be: If a classical device can pass either challenge, then the device can be rewound to hold *both* a valid equation and a preimage, and we will show that knowledge of both is sufficient to break the TCF. Since quantum measurements cannot be rewound this argument does

not apply to quantum machines; if a quantum machine passes the preimage test, it cannot then be used to pass an equation test (and vice versa), since the measurement would cause its state to collapse.

Showing that knowledge of both a preimage and an equation is sufficient to break the TCF presents a new challenge. Specifically, we wish to claim that no efficient (classical or quantum) algorithm can produce both a preimage x_b , as well as an n -string d and a bit c such that $c = d \cdot (x_0 \oplus x_1)$ (even with probability $1/2 + \epsilon$). This may be thought of as a hardcore bit property for the TCF, for the bit of the n -bit string $(x_0 \oplus x_1)$ specified by d . The difficulty is that the specification of the hardcore bit d can be chosen by the quantum device after it gets to see the particular TCF chosen by the verifier, as well as the image y . In this sense, what is required is establishing that the TCF has a kind of “adaptive hardcore bit property.” We describe this property next.

The adaptive hardcore bit property. The adaptive hardcore bit described above is a crucial ingredient in classically testing quantum computers, yet it has not been studied in classical cryptography. Luckily, it turns out that it can be built by relying on structural properties of the well-studied **learning with errors (LWE)** assumption; more specifically, it relies on a property called *leakage resilience*. In this section, we give an overview of the ideas required to prove the adaptive hardcore bit. We begin by describing how the learning with errors assumption can be used to construct a trapdoor claw-free function pair, and then describe how the leakage resilience properties of LWE imply the validity of the adaptive hardcore bit for this construction.

Recall that the learning with errors problem starts with a system of m' linear equations modulo q on a set of n' variables, with $m' > n'$. Starting with a uniformly random matrix $A \in \mathbb{Z}_q^{m' \times n'}$, and a vector $s \in \mathbb{Z}_q^{n'}$ and letting $t = As$ results in an easily solvable linear system of equations (A, t) . To make the inversion problem challenging, a noise vector $e \in \mathbb{Z}_q^{m'}$ is added, so instead $t = As + e$. The distribution over the noise vector e is judiciously chosen (from a suitable Gaussian distribution) so while s is uniquely determined by (A, t) , it is computationally difficult to recover it. The learning with errors assumption states that the distribution over (A, t) is computationally indistinguishable from the distribution over (A, u) for a uniformly random string $u \in \mathbb{Z}_q^{m'}$; in other words, the addition of the noise e computationally hides s .

Given an LWE sample $(A, t = As + e)$, it is natural to try to define a TCF family by letting $f_0(x) = Ax + e_0$ and $f_1(x) = Ax + e_0 + t$. Note that the output of each function is now a random sample from a distribution, since e_0 is randomly chosen. Substituting $t = As + e$, we see that $f_1(x) = A(x + s) + e_0 + e$. If e were 0, then this would mean that $f_1(x) = f_0(x + s)$ (i.e., the two distributions are the same). By sampling e_0 from a Gaussian much wider than e , we can ensure that the distributions $f_1(x)$ and $f_0(x + s)$ are statistically close, thus effectively ensuring that $f_1(x) = f_0(x + s)$. We refer to such a function pair as a noisy trapdoor claw-free function pair (NTCF). Each claw of such a function pair will now have the following property: for all claws (x_0, x_1, y) of the function pair, $x_1 = x_0 - s$. Note that the claw-free property of this pair of functions follows immediately from the LWE assumption, since knowledge of both x_0 and x_1 reveals the secret s .

A quantum device can use an NTCF to set up a superposition over a claw: $\frac{1}{\sqrt{2}}(|0, x_0\rangle + |1, x_0 - s\rangle)$. This follows easily by observing that it can create the superposition $\sum_b \sum_x \sum_{e_0} |b\rangle |x\rangle |Ax + e_0 + bt\rangle$ (omitting normalization factors), and measure the last register to obtain y , creating the desired superposition in the first two registers. Recall that in our earlier description the quantum device worked over qubits, whereas we worked modulo q while defining the NTCF. This is easily remedied by converting all mod q entries to binary strings – letting $n = n' \lceil \log q \rceil$ and $m = m' \lceil \log q \rceil$, we may think of $f_0, f_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$. It is worth remarking at this point that given the form of the superposition over the claw $\frac{1}{\sqrt{2}}(|0, x_0\rangle + |1, x_0 - s\rangle)$, it is tempting to ask why one cannot apply

standard period-finding quantum algorithms to compute s . The point is that even though x_0 and $x_0 - s$ are stored in binary, $x_0 - s$ is computed modulo q and is incompatible with Fourier sampling performed modulo 2. As we will see shortly, this mixing of $\mathbb{Z}_q$ vectors with the Fourier transform mod 2 is what makes the proof of the adaptive hardcore bit possible.

Given this additional structure, we can now state the adaptive hardcore bit property a bit more precisely. The adaptive hardcore bit property states that it is difficult to hold both a single preimage x_b , as well as a string $d \in \{0, 1\}^n \setminus 0^n$ and a bit c such that $c = d \cdot (x_0 \oplus x_1)$. Since $x_1 = x_0 - s$, one might hope to express the last condition directly in terms of s . Note that $d \cdot (x_0 \oplus (x_0 - s))$ is *not* equal to $d \cdot s$, due to the fact that a binary XOR and a difference modulo q do not cancel. Instead, if s is binary, and can therefore be interpreted as a string in $\{0, 1\}^{n'}$, then it is possible to use x_b to efficiently compute a string $d' \in \{0, 1\}^{n'}$ such that $d \cdot (x_0 \oplus (x_0 - s)) = d' \cdot s$, via a linear map that relies only on the fact that s is binary. To see how this map works, consider the special case in which $n' = 1$ and the secret s is a single bit. Then, if we let $d' = d \cdot (x_0 \oplus (x_0 - 1))$ ($x_0 - 1$ is again performed modulo q), then it follows that $d \cdot (x_0 \oplus (x_0 - s)) = d' \cdot s$. This reasoning immediately extends to the general case. It follows that the adaptive hardcore bit can be reformulated as stating that it is difficult to produce a string d' and a bit c such that $d' \cdot s = c$; in other words, not only is it hard to find the secret s , it is even hard to find any bit of the secret s .

If we could assume that d' was chosen independently of the LWE sample, then the desired hardcore bit property would follow from current results in classical cryptography that strengthen the security of LWE to prove leakage resilience: Given an LWE sample $As + e$, any given bit of s is computationally indistinguishable from a uniformly random bit. Unfortunately, there is an added difficulty in our setting: The quantum device can choose the string d' after seeing the LWE sample (after all, the device requires the LWE sample $t = As + e$ to evaluate the function). It is in this sense that the hardcore bit property is adaptive. We now outline the leakage resilience argument to describe how it can be adapted to our setting.

In proving leakage resilience [21], the matrix A is replaced with a computationally indistinguishable matrix $BC + E$, where $C \in \mathbb{Z}_q^{\ell' \times n'}$, for $\ell' \ll n'$. The computational indistinguishability is immediately implied by treating $BC + E$ as a smaller LWE sample, in which C is the secret. Moreover, E is chosen from a Gaussian with width sufficiently smaller than the Gaussian noise e , implying that $(BC + E)s + e$ is statistically close to $BCs + e$. The point of this substitution is that the matrix C compresses s , and the leftover hash lemma can be invoked to argue that even given Cs (which is at least as much information as $(BC + E)s$), any bit of s is statistically close to uniform, thus showing that it is a hardcore bit given A .

In the situation we are interested in, the choice of d' may depend upon the LWE sample $As + e$, which corresponds in the leakage resilience argument to d' depending on Cs . We wish to argue that $d' \cdot s$ is still statistically close to uniform. This is where the mod q versus mod 2 difference comes into play: In our setting, the string d' is binary (as is the inner product of d' and s), whereas the entries of C are uniformly random entries in $\mathbb{Z}_q$. It can be shown via a Fourier analytic argument that even if Cs is fixed, then there is enough entropy left in s that $d' \cdot s$ is statistically close to uniform.

As you might expect, the argument outlined above requires that d' is non-zero. Therefore, the verifier must check that the string d returned by the prover yields a non-zero d' . Note that although d' is an easily computable function of x_b and d , the verifier has no way of knowing which preimage (x_0 or x_1) the adversary may have in mind. It follows that there are two different values d'_0, d'_1 (corresponding to x_0 and x_1) that *both* must be checked to be non-zero. Of course, the verifier knows the trapdoor and can perform this check efficiently.

The problem is a little more serious in the adaptive hardcore bit proof, which requires that the validity of d' be testable efficiently without the trapdoor to maintain the entropy of s . Clearly, knowledge of both preimages x_0 and x_1 is not the answer, since this uniquely determines s . Instead, we modify the protocol so the verifier imposes a more restrictive constraint on d , by only accepting d such that the first half of d'_0 is non-zero, and the second half of d'_1 is non-zero. Observe that checking whether d' satisfies these constraints can either be done with x_0 and the second half of s (combined with x_0 , this can be used to compute the second half of x_1 , and therefore the second half of d'_1), or x_1 and the first half of s . Moreover, observe that with this limited knowledge, the adaptive hardcore bit still holds; for example, knowledge of only x_0 and the second half of s preserves the entropy of the first half of s , thereby allowing us to apply the hardcore bit argument to the first half of d'_0 (which we know to be non-zero).

Quantum Supremacy. The qubit certification protocol described above has implications for an important milestone in the experimental realization of quantum computers, namely “quantum supremacy”: a proof that an (untrusted) quantum computing device performs some computational task that cannot be solved classically without impractical resources. While this could in principle be achieved by demonstrating quantum factoring, the latter requires quantum resources well beyond the capability of near-term experiments. Instead current proposals are based on sampling problems (see, e.g., Reference [24] for a recent survey). The major challenge for these proposals is verifying that the quantum computer did indeed sample from the desired probability distribution, and all existing proposals rely on exponential time classical algorithms for verification. By contrast, our supremacy provides a proof of quantumness that can be verified by a classical verifier in polynomial time. This proposal seems promising from a practical viewpoint—indeed, even using off-the-shelf bounds for LWE-based cryptography suggests that a protocol providing 50 bits of security could be implemented with a quantum device of around 2,000 qubits (see, e.g., Reference [25]). It would be worth exploring whether there are clever implementations of this scheme that can lead to a protocol in the 200–500 qubit range.

Another challenge in making our proposal suitable for near-term devices is fault tolerance. While our protocol will require some level of fault tolerance, the hope is that it might not require general fault tolerance techniques, due to its robustness: Our protocol is robust to a device that only successfully answers the verifier’s challenges with a sufficiently large, but constant, success probability.

Certifiable randomness. The challenge in achieving certifiable randomness lies in using computational assumptions to establish not pseudorandomness, but rather that the output of the protocol must be (close to) statistically random. In our analysis, we leverage the properties of the NTCF to characterize the quantum state and measurements of the untrusted quantum device—essentially showing that it must have a qubit initialized in state $|+\rangle$, which it measured in the standard basis, thus generating one bit of statistical randomness. This is the analogue of the use of the violation of Bell inequalities to characterize the state of the device in entanglement-based testing.

We first explain how to show that a device that succeeds in the qubit certification test (which we will often refer to as a single round test or single round protocol) must generate randomness. In the test the device must make one of two measurements: either a “preimage” measurement or an “equation” measurement. We focus on a single bit of information provided by each measurement. The “preimage” measurement can be treated as a projection into one of two orthogonal subspaces corresponding to the two preimages x_0, x_1 for the element y that the device has returned to the verifier. The “equation” measurement can similarly be coarse-grained into a projection on one of two orthogonal subspaces, “valid” or “invalid,” i.e., the subspace that corresponds to all measure-

ment outcomes d, c such that $d \cdot (x_0 + x_1) = c$, or the subspace associated with outcomes d, c such that $d \cdot (x_0 + x_1) = c \oplus 1$.

Applying Jordan’s lemma, it is possible to decompose the device’s Hilbert space into a direct sum of one- and two-dimensional subspaces, such that within each two-dimensional subspace the “preimage” and “equation” measurements each correspond to an orthonormal basis, such that the two bases make a certain angle with each other. We argue that almost all angles must be very close to $\pi/4$. Indeed, whenever the angles are *not* near-maximally unbiased, it is possible to show that by considering the effect of performing the measurements in sequence, one can devise an “attack” on the NTCF of a kind that contradicts the adaptive hardcore bit property of the NTCF—informally, the attack can simultaneously produce a valid preimage and a valid equation, with non-negligible advantage.

As a result, it is possible to show that the state and (coarse-grained) measurements of the device are, up to a global change of basis, close to the following: The device starts with a qubit initialized to $|+\rangle$, which it measures in the standard basis for the case of a preimage test and in the Hadamard basis for the case of an equation test. The fact that an efficient quantum device cannot break the cryptographic assumption has thus been translated into a characterization of the state and actions of the quantum device, which further implies that the output of the device in the single round test must contain close to a bit of true (information theoretic) randomness.

One might further conjecture that for a generic TCF (e.g., modeled as a random oracle), if the output of any efficient quantum device passes the single round test with non-negligible advantage over $\frac{1}{2}$, then the triple y, d, c returned in the equation test must have high min-entropy. Such a strong statement would immediately yield a randomness certification protocol. Among the many difficulties in showing such a statement is that both y and d may be adaptively and adversarially chosen—in the single round protocol above this issue is addressed by the adaptive hardcore bit property of the NTCF.

Outline of randomness generation protocol. Going beyond the analysis of the single round test requires significantly more work. So far, we have argued that if an efficient quantum algorithm has the ability to generate a valid equation with probability sufficiently close to 1, then, if instead it is asked for a preimage, this preimage must be close to uniformly distributed over the two possibilities. To leverage this our randomness expansion protocol proceeds in multiple rounds, repeatedly asking for new images y and a preimage of y (to generate randomness) while inserting a few randomly located equation tests to test the device. Each time an “equation” challenge has been answered, we refresh the pseudorandom keys used for the NTCF. This is required to avoid a simple “attack” by the device, which would repeatedly use the same y , preimage x , and guessed equation d —succeeding in the protocol with probability $\frac{1}{2}$ without generating any randomness.

Let us call the sequence of rounds with a particular set of pseudorandom keys an epoch. Intuitively, we would like to claim that if the device passes all the equation tests, then for most epochs and for most rounds within that epoch, the state of the device and its measurements must be (close to) as characterized above: It starts with a qubit initialized to $|+\rangle$, which it measures in the standard basis for the case of a preimage test, and in the Hadamard basis for the case of an equation test. To show this, we would like to claim that if the device passes all the equation tests, then for most such tests it must produce a valid equation with probability close to 1. Since each equation test occurs at a random round in the epoch, it should follow from the adaptive hardcore bit property that the sequence of bits that the verifier extracts from the device’s answers to preimage tests during that epoch must look statistically random. We give a martingale-based argument to formalize this intuition.

There is, however, a bigger challenge to analyzing the protocol—we must show that the sequence that the verifier extracts from the device’s answers to preimage tests must look statistically random even to an infinitely powerful quantum adversary, who may share an arbitrary entangled state with the quantum device. If we could assert that each round of the protocol is played with a qubit exactly in state $|+\rangle$, and measured in the standard basis for the case of a preimage test, then this would lead to an easy proof that the extracted sequence looks random to the adversary. Unfortunately, the characterization of the device’s qubits leaves plenty of room for entanglement with the adversary. Showing that such entanglement cannot leak too much information about the device’s measurements was the major challenge in previous work on certified randomness through Bell inequality violations [6, 30, 43]. Our cryptographic setting presents a new difficulty, which is that in contrast to the Bell inequality violation scenarios, in our setting it is not *impossible* for a deterministic device to succeed in the test: It is merely *computationally hard* to do so. This prevents us from directly applying the results in References [6, 30] and requires us to suitably modify their framework. We describe this part of the argument in more detail below.

In terms of efficiency, for the specific LWE-based NTCF that we construct, our protocol can use as few as $\text{poly}(\log(N))$ bits of randomness to generate $O(N)$ bits that are statistically within negligible distance from uniform. However, this requires assuming that the underlying LWE assumption is hard even for sub-exponential size quantum circuits with polynomial-size quantum advice (which is consistent with current knowledge). The more conservative assumption that our variant of LWE is only hard for polynomial size quantum circuits requires $O(N^\epsilon)$ bits of randomness for generating the NTCF, for any constant $\epsilon > 0$. The following is an informal description; see Theorem 8.13 for a more formal statement.

THEOREM 1.1 (INFORMAL). *Let $\mathcal{F}$ be an NTCF family and λ a security parameter. Let $N = \Omega(\lambda^2)$ and assume the quantum hardness of solving lattice problems of dimension λ in time $\text{poly}(N)$. There is an N -round protocol for the interaction between a classical polynomial-time verifier and a quantum polynomial-time device such that the protocol can be executed using $\text{poly}(\log(N), \lambda)$ bits of randomness, and for any efficient device and side information E correlated with the device’s initial state,*

$$H_{\infty}^{\delta}(O|CE)_{\bar{p}} \geq (\xi - o(1))N.$$

Here, ξ is a positive constant, δ is a negligible function of λ , and $\bar{p}$ is the final state of the classical output register O , the classical register C containing the verifier’s messages to the device, and the side information E , restricted to transcripts that are accepted by the verifier in the protocol.

Sketch of the security analysis. We describe the protocol in slightly more detail (see Section 5 for a formal description). The verifier first uses $\text{poly}(\log(N), \lambda)$ bits of randomness to select a pair of functions $\{f_{k,b}\}_{b \in \{0,1\}}$ from an NTCF family and sends the public function key k to the quantum device. This pair of functions can be interpreted as a single 2-to-1 function $f_k : (b, x) \mapsto f_{k,b}(x)$. The verifier keeps private the trapdoor information that allows to invert f_k . The protocol then proceeds for N rounds. In each round the device first outputs a value y in the common range of $f_{k,0}$ and $f_{k,1}$. After having received y , the verifier issues one of two challenges: 0 or 1, preimage or equation. If the challenge is “preimage,” then the device must output an x such that $f(x) = y$. If the challenge is “equation,” then the device must output a nontrivial binary vector d and a bit c such that $d \cdot (x_0 + x_1) = c$, where x_0 and x_1 are the unique preimages of y under $f_{k,0}$ and $f_{k,1}$, respectively. Since the verifier has the secret key, she can efficiently compute x_0 and x_1 from y and therefore check the correctness of the device’s response to each challenge. The verifier chooses $\text{poly}(\log(N))$ rounds in which to issue the challenge 1, or “equation,” at random. Selecting these rounds requires only $\text{poly}(\log(N))$ random bits. At the end of each such round, the verifier samples a new pair of

functions from the NTCF family and communicates the new public key to the device. On each of the remaining $N - \text{poly} \log(N)$ rounds the verifier records a bit according to whether the device returns the preimage x_0 , or x_1 (e.g., recording 0 for the lexicographically smaller preimage). At the end of the protocol the verifier uses a strong quantum-proof randomness extractor to extract $\Omega(N)$ bits of randomness from the recorded string (this requires at most an additional $\text{poly} \log(N)$ bits of uniformly random seed).

To guarantee that the extractor produces bits that are statistically close to uniform, we would like to prove that the $N - \text{poly} \log(N)$ random bits recorded by the verifier must have $\Omega(N)$ bits of (smoothed) min-entropy,¹ even conditioned on the side information available to an infinitely powerful quantum adversary, who may share an arbitrary entangled state with the quantum device.

The analysis proceeds as follows: First, we assume without loss of generality that the entire protocol is run coherently, i.e., we may assume that the initial state of the quantum device (holding quantum register D) and the adversary (holding quantum register E) is a pure state $|\phi\rangle_{DE}$, since the adversary may as well start with a purification of their joint state. We may also assume that the verifier starts with a cat state on $\text{poly} \log(N)$ qubits and uses one of the registers of the state, C, to provide the random bits used to select the type of test being performed in each round. (This is for the sake of analysis only; the actual verifier is of course completely classical.) We can similarly arrange that the state remains pure throughout the protocol by using the principle of deferred measurement. Our goal is to show a lower bound on the smooth min-entropy of the output register O in which the verifier has recorded the device's outputs, conditioned on the state E of the adversary, and on the register C of the cat state (conditioning on the latter represents the fact that the verifier's choice of challenges may be leaked to the adversary, and we would like security even in this scenario). Intuitively, this amounts to bounding the information accessible to the most powerful adversary quantum mechanics allows, conditioned on the joint state of the verifier and device.

To bound the entropy of the final state, we need to show that the entropy "accumulates" at each round of the protocol. A general framework to establish entropy accumulation in quantum protocols such as the one considered here was introduced in Reference [6]. At a high level, the approach consists in reducing the goal of a min-entropy bound to a bound on the appropriate notion of $(1 + \epsilon)$ quantum conditional Rényi entropy and then arguing that, under suitable conditions on the process that generates the outcomes recorded in the protocol, entropy accumulates sequentially throughout the protocol.

In a little more detail, the first step on getting a handle on the smooth min-entropy is to use the **quantum asymptotic equipartition property (QAEF)** [42] to relate it to the $(1 + \epsilon)$ conditional Rényi entropy for suitably small ϵ . The second step uses a duality relation for the conditional Rényi entropy to relate the $(1 + \epsilon)$ conditional Rényi entropy of the output register O, conditioned on the adversary side information in R and the register C of the cat state, to a quantity analogous to the $(1 - \epsilon')$ conditional Rényi entropy of the output register, conditioned on the register E for the device, and a purifying copy of the register C of the cat state. The latter quantity, a suitable conditional entropy of the output register conditioned on the challenge register and the state of the device, is the quantity that we ultimately aim to bound. Note what these transformations have achieved for us: It is now sufficient to consider as side information only "known" quantities in the protocol, the verifier's choice of challenges and the device's state; the information held by the adversary plays no other role than that of a purifying register.

As mentioned earlier, our cryptographic setting presents the additional difficulty that our guarantee is only that it is computationally hard for a deterministic device to succeed in the protocol. The results in References [6, 30] crucially rely on the fact that the process that generates the ran-

¹We refer to Section 2 for definitions of entropic quantities.

domness does so irrespective of the quantum state in which it is initialized (as long as the output of the process satisfies the test’s success criterion). This requirement comes from the conditioning that is performed to show that entropy accumulates; in our setting, conditioning is more delicate, as it can in principle induce non-computationally efficient states for the device.

Recall that we argued that for a single round of the protocol, we can decompose the device’s Hilbert space into a direct sum of one- or two-dimensional subspaces, such that within most two-dimensional subspace the “preimage” and “equation” measurements correspond to orthonormal bases that make an angle close to $\pi/4$ with each other. Showing that the Rényi entropy accumulates in each round requires a device in which *all* angles are close to $\pi/4$, not “almost all.” To accommodate for this, we “split” the state of the device into its component on the good subspace, where the angles are unbiased, and the bad subspace, where the measurements may be aligned. The fact that the distinction between good and bad subspace is not measured in the protocol, but is only a distinction made for the analysis, requires us to apply a fairly delicate martingale based argument that takes into account possible interference effects and bounds those “branches” where the state has gone through the bad subspace an improbably large number of times. Whenever the state lies in the good subspace, we can appeal to an uncertainty principle from Reference [30] to show that the device’s measurement increases the conditional Rényi entropy of the output register by a small additive constant. Pursuing this approach across all N rounds, we obtain a linear lower bound on the conditional Rényi entropy of the output register, conditioned on the state of the device. As argued above this in turn translates into a linear lower bound on the smooth conditional min-entropy of the output, conditioned on the state of the adversary and the verifier’s choice of challenges. It only remains to apply a quantum-proof randomness extractor to the output, using a poly-logarithmic number of additional bits of randomness, to obtain the final result.

Concurrent and related work. The idea of using a TCF as a basic primitive in interactions between an efficient quantum prover and a classical verifier has been further developed in recent work by Mahadev [27], giving the first construction of a quantum fully homomorphic encryption scheme with classical keys. In further follow-up work, Mahadev [28] shows a remarkable use of an NTCF family with adaptive hardcore bit. Namely, that the NTCF can be used to certify that a prover measures a qubit in a prescribed basis (standard or Hadamard). This allows to achieve single prover *verifiability* for quantum computations using a purely classical verifier (but relying on computational assumptions).

Independently of this work, a construction of trapdoor one-way functions with second preimage resistance based on LWE was recently introduced in Reference [14], where it is used to achieve delegated computation in the weaker honest-but-curious model for the adversary (i.e., without soundness against provers not following the protocol). The family of functions considered in Reference [14] is not sufficient for our purposes, as it lacks the adaptive hardcore bit property.

After the completion of our work, in Reference [20] the construction of NTCF family introduced here was extended to a more general hardcore bit property (informally, over $\mathbb{Z}_8$ instead of $\mathbb{Z}_2$ here) and used to implement a two-party functionality called “remote state preparation” by which a classical client can “force” the preparation of one out of eight possible single-qubit quantum states by the prover. The authors of Reference [15] also generalize Reference [14] to obtain a similar functionality; however, their construction does not offer the property of being verifiable (informally, it is possible for the server to prepare a state that is not the expected one).

We believe that the technique of constraining the power of a quantum device using NTCFs promises to be a powerful tool for the field of untrusted quantum devices.

Organization. We start with some notation and preliminaries in Section 2. Section 3 contains the definition of a **noisy trapdoor claw-free family (NTCF)**. Our construction for such a family

is given in Section 4 (with Appendix 2.3 containing relevant preliminaries on the learning with errors problem). The randomness generation protocol is described in Section 5. In Section 6, we introduce our formalism for modeling the actions of an arbitrary prover, or device, in the protocol. In Section 7, we analyze a single round of the protocol, and in Section 8, we show that randomness accumulates across multiple rounds.

2 PRELIMINARIES

2.1 Notation

$\mathbb{Z}$ is the set of integers, and $\mathbb{N}$ the set of natural numbers. For any $q \in \mathbb{N}$ such that $q \geq 2$, we let $\mathbb{Z}_q$ denote the ring of integers modulo q . We generally identify an element $x \in \mathbb{Z}_q$ with its unique representative $[x]_q \in (-\frac{q}{2}, \frac{q}{2}] \cap \mathbb{Z}$. For $x \in \mathbb{Z}_q$, we define $|x| = |[x]_q|$. When considering an $s \in \{0, 1\}^n$, we sometimes also think of s as an element of $\mathbb{Z}_q^n$, in which case, we write it as s .

We use the terminology of polynomially bounded and negligible functions. A function $n : \mathbb{N} \rightarrow \mathbb{R}_+$ is *polynomially bounded* if there exists a polynomial p such that $n(\lambda) \leq p(\lambda)$ for all $\lambda \in \mathbb{N}$. A function $n : \mathbb{N} \rightarrow \mathbb{R}_+$ is *negligible* if for every polynomial p , $p(\lambda)n(\lambda) \rightarrow_{\lambda \rightarrow \infty} 0$. We write $\text{negl}(\lambda)$ to denote an arbitrary negligible function of λ . For two parameters κ, λ , we write $\kappa \ll \lambda$ to express the constraint that κ should be “sufficiently smaller than” λ , meaning that there exists a small universal constant $c > 0$ such that $\kappa \leq c\lambda$, where c is usually implicit for context.

$\mathcal{H}$ always denotes a finite-dimensional Hilbert space. We use indices $\mathcal{H}_A, \mathcal{H}_B$, and so on, to refer to distinct spaces. $\text{Pos}(\mathcal{H})$ is the set of positive semidefinite operators on $\mathcal{H}$, and $\text{D}(\mathcal{H})$ the set of density matrices, i.e., the positive semidefinite operators with trace 1. For an operator X on $\mathcal{H}$, we use $\|X\|$ to denote the operator norm (largest singular value) of X , and $\|X\|_{tr} = \frac{1}{2}\|X\|_1 = \frac{1}{2}\text{Tr}\sqrt{XX^\dagger}$ for the trace norm.

2.2 Distributions

We generally use the letter D to denote a distribution over a finite domain X , and f for a density on X , i.e., a function $f : X \rightarrow [0, 1]$ such that $\sum_{x \in X} f(x) = 1$. We often use the distribution and its density interchangeably. We write U for the uniform distribution. We write $x \leftarrow D$ to indicate that x is sampled from distribution D , and $x \leftarrow_U X$ to indicate that x is sampled uniformly from the set X . We write $\mathcal{D}_X$ for the set of all densities on X . For any $f \in \mathcal{D}_X$, $\text{SUPP}(f)$ denotes the support of f ,

$$\text{SUPP}(f) = \{x \in X \mid f(x) > 0\}.$$

For two densities f_1 and f_2 over the same finite domain X , the Hellinger distance between f_1 and f_2 is

$$H^2(f_1, f_2) = 1 - \sum_{x \in X} \sqrt{f_1(x)f_2(x)}. \quad (3)$$

The total variation distance between f_1 and f_2 is

$$\|f_1 - f_2\|_{TV} = \frac{1}{2} \sum_{x \in X} |f_1(x) - f_2(x)| \leq \sqrt{2H^2(f_1, f_2)}. \quad (4)$$

The following immediate lemma relates the Hellinger distance and the trace distance of superpositions:

LEMMA 2.1. *Let X be a finite set and $f_1, f_2 \in \mathcal{D}_X$. Let*

$$|\psi_1\rangle = \sum_{x \in X} \sqrt{f_1(x)}|x\rangle \quad \text{and} \quad |\psi_2\rangle = \sum_{x \in X} \sqrt{f_2(x)}|x\rangle.$$

Then

$$\| |\psi_1\rangle - |\psi_2\rangle \|_{tr} = \sqrt{1 - (1 - H^2(f_1, f_2))^2}.$$

We say that a family of quantum circuits $\{C_\lambda\}_{\lambda \in \mathbb{N}}$ (respectively, observables $\{O_\lambda\}_{\lambda \in \mathbb{N}}$) is *polynomial-time generated* if there exists a polynomial-time deterministic Turing machine that, on every input $\lambda \in \mathbb{N}$, returns a gate-by-gate encoding of the circuit C_λ (respectively, of a circuit that implements O_λ). We introduce a notion of efficient distinguishability between distributions.

Definition 2.2. We say that two families of distributions $D_0 = \{D_{0,\lambda}\}_{\lambda \in \mathbb{N}}$ and $D_1 = \{D_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ on the same finite set $\{X_\lambda\}$ are *computationally indistinguishable* if for every polynomial-time generated family of quantum circuits $\mathcal{A} = \{A_\lambda : X_\lambda \rightarrow \{0, 1\}\}$ it holds that

$$\left| \Pr_{x \leftarrow D_{0,\lambda}} [A_\lambda(x) = 0] - \Pr_{x \leftarrow D_{1,\lambda}} [A_\lambda(x) = 0] \right| = \text{negl}(\lambda), \quad (5)$$

where the probability is taken over the choice of x from either distribution as well as randomness inherent in any measurement performed by the circuit A_λ .

The next definition generalizes the previous one to the case of quantum states.

Definition 2.3. We say that two families of sub-normalized density matrices $\sigma_0 = \{\sigma_{0,\lambda}\}_{\lambda \in \mathbb{N}}$ and $\sigma_1 = \{\sigma_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ on the same Hilbert space $\{\mathcal{H}_\lambda\}$ are *computationally indistinguishable* if for every polynomial-time generated family of observables $O = \{O_\lambda\}_{\lambda \in \mathbb{N}}$ it holds that

$$\left| \text{Tr}(O_\lambda(\sigma_{0,\lambda} - \sigma_{1,\lambda})) \right| = \text{negl}(\lambda).$$

2.3 The Learning with Errors Problem

We give some background on the **Learning with Errors problem (LWE)**. For a positive real B and a positive integer q , the truncated discrete Gaussian distribution over $\mathbb{Z}_q$ with parameter B is the distribution supported on $\{x \in \mathbb{Z}_q : \|x\| \leq B\}$ with density

$$D_{\mathbb{Z}_q, B}(x) = \frac{e^{-\frac{\pi \|x\|^2}{B^2}}}{\sum_{x \in \mathbb{Z}_q, \|x\| \leq B} e^{-\frac{\pi \|x\|^2}{B^2}}}. \quad (6)$$

More generally, for a positive integer m the truncated discrete Gaussian distribution over $\mathbb{Z}_q^m$ with parameter B is the distribution supported on $\{x \in \mathbb{Z}_q^m : \|x\| \leq B\sqrt{m}\}$ with density

$$\forall x = (x_1, \dots, x_m) \in \mathbb{Z}_q^m, \quad D_{\mathbb{Z}_q^m, B}(x) = D_{\mathbb{Z}_q, B}(x_1) \cdots D_{\mathbb{Z}_q, B}(x_m). \quad (7)$$

LEMMA 2.4. Let B be a positive real and q, m positive integers. Consider $\mathbf{e} \in \mathbb{Z}_q^m$ such that $\|\mathbf{e}\| \leq B\sqrt{m}$. The Hellinger distance between the distribution $D = D_{\mathbb{Z}_q^m, B}$ and the shifted distribution $D + \mathbf{e}$, with density $(D + \mathbf{e})(x) = D(x - \mathbf{e})$, satisfies

$$H^2(D, D + \mathbf{e}) \leq 1 - e^{-\frac{2\pi\sqrt{m}\|\mathbf{e}\|}{B}}, \quad (8)$$

and the statistical distance between the two distributions satisfies

$$\|D - (D + \mathbf{e})\|_{TV}^2 \leq 2 \left(1 - e^{-\frac{2\pi\sqrt{m}\|\mathbf{e}\|}{B}} \right). \quad (9)$$

PROOF. Let $\tau = \sum_{x \in \mathbb{Z}_q, \|x\| \leq B} e^{-\frac{\pi \|x\|^2}{B^2}}$. We will rely on the fact that for any $\mathbf{e}_0$ in the support of $D_{\mathbb{Z}_q^m, B}$, $\|\mathbf{e}_0\| \leq B\sqrt{m}$. We can compute the bound as follows:

$$\begin{aligned}
 \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, B}(\mathbf{e}_0) D_{\mathbb{Z}_q^m, B}(\mathbf{e}_0 - \mathbf{e})} &= \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi (\|\mathbf{e}_0\|^2 + \|\mathbf{e}_0 - \mathbf{e}\|^2)}{2B^2}} \\
 &\geq \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi (\|\mathbf{e}_0\|^2 + (\|\mathbf{e}_0\| + \|\mathbf{e}\|)^2)}{2B^2}} \\
 &= \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi (\|\mathbf{e}_0\|^2)}{B^2}} e^{-\frac{\pi (2\|\mathbf{e}_0\|\|\mathbf{e}\|)}{2B^2}} e^{-\frac{\pi (\|\mathbf{e}\|^2)}{2B^2}} \\
 &\geq e^{-\frac{\pi (\|\mathbf{e}\|^2 + 2B\sqrt{m}\|\mathbf{e}\|)}{2B^2}} \frac{1}{\tau^m} \sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} e^{-\frac{\pi (\|\mathbf{e}_0\|)^2}{B^2}} \\
 &= e^{-\frac{\pi (\|\mathbf{e}\|^2 + 2B\sqrt{m}\|\mathbf{e}\|)}{2B^2}} \\
 &\geq e^{-\frac{\pi (4B\sqrt{m}\|\mathbf{e}\|)}{2B^2}} \\
 &= e^{-\frac{2\pi \sqrt{m}\|\mathbf{e}\|}{B}}.
 \end{aligned}$$

The bound on the statistical distance follows from the bound on the Hellinger distance using the inequality in (4). $\square$

We define the main assumption that underlies all computational hardness claims made in the article.

Definition 2.5. For a security parameter λ , let $n, m, q \in \mathbb{N}$ be integer functions of λ . Let $\chi = \chi(\lambda)$ be a distribution over $\mathbb{Z}$. The $\text{LWE}_{n, m, q, \chi}$ problem is to distinguish between the distributions $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q})$ and $(\mathbf{A}, \mathbf{u})$, where $\mathbf{A} \leftarrow_U \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow_U \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{u} \leftarrow_U \mathbb{Z}_q^m$. Often, we consider the hardness of solving LWE for any function m such that m is at most a polynomial in $n \log q$. This problem is denoted $\text{LWE}_{n, q, \chi}$.

In this article, we make the assumption that no quantum polynomial-time procedure can solve the $\text{LWE}_{n, q, \chi}$ problem with more than a negligible advantage in λ , even when given access to a quantum polynomial-size advice state depending on the parameters n, m, q , and χ of the problem. We refer to this assumption as “the $\text{LWE}_{n, q, \chi}$ assumption.”

As shown in References [34, 37], for any $\alpha > 0$ such that $\sigma = \alpha q \geq 2\sqrt{n}$ the $\text{LWE}_{n, q, D_{\mathbb{Z}_q, \sigma}}$ problem, where $D_{\mathbb{Z}_q, \sigma}$ is the discrete Gaussian distribution, is at least as hard as approximating the **shortest independent vector problem** (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$, where $\tilde{O}$ hides factors logarithmic in the argument, in *worst case* dimension n lattices. This is proven using a quantum reduction. Classical reductions (to a slightly different problem) exist as well [11, 33] but with somewhat worse parameters. The best known (classical or quantum) algorithm for these problems runs in time $2^{\tilde{O}(n/\log \gamma)}$. For our construction, given in Section 4, we assume hardness of the problem against a quantum polynomial-time adversary in the case that γ is a super polynomial function in n . This is a commonly used assumption in cryptography (for, e.g., homomorphic encryption schemes such as Reference [19]).

We use two additional properties of the LWE problem. The first is that it is possible to generate LWE samples $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ such that there is a trapdoor allowing recovery of $\mathbf{s}$ from the samples.

THEOREM 2.6 (THEOREM 5.1 IN [29]). *Let $n, m \geq 1$ and $q \geq 2$ be such that $m = \Omega(n \log q)$. There is an efficient randomized algorithm $\text{GENTRAP}(1^n, 1^m, q)$ that returns a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a trapdoor $t_{\mathbf{A}}$ such that the distribution of $\mathbf{A}$ is negligibly (in n) close to the uniform distribution. Moreover, there is an efficient algorithm INVERT that, on input $\mathbf{A}, t_{\mathbf{A}}$ and $\mathbf{A}s + \mathbf{e}$ where $\|\mathbf{e}\| \leq q/(C_T \sqrt{n \log q})$ and C_T is a universal constant, returns s and $\mathbf{e}$ with overwhelming probability over $(\mathbf{A}, t_{\mathbf{A}}) \leftarrow \text{GENTRAP}(1^n, 1^m, q)$.*

The second property is the existence of a “lossy mode” for LWE. The following definition is Definition 3.1 in Reference [5]:

Definition 2.7. Let $\chi = \chi(\lambda)$ be an efficiently sampleable distribution over $\mathbb{Z}_q$. Define a lossy sampler $\tilde{\mathbf{A}} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, \chi)$ by $\tilde{\mathbf{A}} = \mathbf{B}\mathbf{C} + \mathbf{F}$, where $\mathbf{B} \leftarrow_U \mathbb{Z}_q^{m \times \ell}$, $\mathbf{C} \leftarrow_U \mathbb{Z}_q^{\ell \times n}$, $\mathbf{F} \leftarrow \chi^{m \times n}$.

THEOREM 2.8 (LEMMA 3.2 IN [5]). *Under the $\text{LWE}_{\ell, q, \chi}$ assumption, the distribution of a random $\tilde{\mathbf{A}} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, \chi)$ is computationally indistinguishable from $\mathbf{A} \leftarrow_U \mathbb{Z}_q^{m \times n}$.*

2.4 Entropies

For $p \in [0, 1]$, we write $H(p) = -p \log p - (1 - p) \log(1 - p)$ for the binary Shannon entropy. We measure randomness using Rényi conditional entropies. For a positive semidefinite matrix $\sigma \in \text{Pos}(\mathcal{H})$ and $\varepsilon \geq 0$, let

$$\langle \sigma \rangle_{1+\varepsilon} = \text{Tr}(\sigma^{1+\varepsilon}).$$

This quantity satisfies the following approximate linearity relations:

$$\forall \varepsilon \in [0, 1], \quad \langle \sigma \rangle_{1+\varepsilon} + \langle \tau \rangle_{1+\varepsilon} \leq \langle \sigma + \tau \rangle_{1+\varepsilon} \leq (1 + O(\varepsilon)) (\langle \sigma \rangle_{1+\varepsilon} + \langle \tau \rangle_{1+\varepsilon}). \quad (10)$$

In addition, for positive semidefinite $\sigma, \rho \in \text{Pos}(\mathcal{H})$ such that the support of ρ is included in the support of σ , and $\varepsilon \geq 0$, let

$$\tilde{Q}_{1+\varepsilon}(\rho \| \sigma) = \langle \sigma^{-\frac{\varepsilon}{2(1+\varepsilon)}} \rho \sigma^{-\frac{\varepsilon}{2(1+\varepsilon)}} \rangle_{1+\varepsilon}. \quad (11)$$

Quantum analogues of the conditional Rényi entropies can be defined as follows:

Definition 2.9. Let $\rho_{\text{AB}} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be positive semidefinite. Given $\varepsilon > 0$, the $(1 + \varepsilon)$ Rényi entropy of A conditioned on B is defined as

$$H_{1+\varepsilon}(A|B)_\rho = \sup_{\sigma \in \text{D}(\mathcal{H}_B)} H_{1+\varepsilon}(A|B)_{\rho|\sigma},$$

where for any $\sigma_B \in \text{D}(\mathcal{H}_B)$,

$$H_{1+\varepsilon}(A|B)_{\rho|\sigma} = -\frac{1}{\varepsilon} \log \tilde{Q}_{1+\varepsilon}(\rho \| \sigma).$$

Rényi entropies are used in the proofs, because they have better “chain-rule-like” properties than the min-entropy, which is the most appropriate measure for randomness quantification.

Definition 2.10. Let $\rho_{\text{AB}} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ be positive semidefinite. Given a density matrix the min-entropy of A conditioned on B is defined as

$$H_\infty(A|B)_\rho = \sup_{\sigma \in \text{D}(\mathcal{H}_B)} H_\infty(A|B)_{\rho|\sigma},$$

where for any $\sigma_B \in \text{D}(\mathcal{H}_B)$,

$$H_\infty(A|B)_{\rho|\sigma} = \max \left\{ \lambda \geq 0 \mid 2^{-\lambda} \text{Id}_A \otimes \sigma_B \geq \rho_{AB} \right\}.$$

It is often convenient to consider the *smooth* min-entropy, which is obtained by maximizing the min-entropy over all positive semidefinite operators matrices in an ε -neighborhood of ρ_{AB} . The definition of neighborhood depends on a choice of metric; the canonical choice is the “purified distance.” Since this choice will not matter for us, we defer to Reference [41] for a precise definition.

Definition 2.11. Let $\varepsilon \geq 0$ and $\rho_{AB} \in \text{Pos}(\mathcal{H}_A \otimes \mathcal{H}_B)$ positive semidefinite. The ε -smooth min-entropy of A conditioned on B is defined as

$$H_{\infty}^{\varepsilon}(A|B)_{\rho} = \sup_{\sigma_{AB} \in \mathcal{B}(\rho_{AB}, \varepsilon)} H_{\infty}(A|B)_{\sigma},$$

where $\mathcal{B}(\rho_{AB}, \varepsilon)$ is the ball of radius ε around ρ_{AB} , taken with respect to the purified distance.

The following theorem relates the min-entropy to the the Rényi entropies introduced earlier. The theorem expresses the fact that, up to a small amount of “smoothing” (the parameter δ in the theorem), all these entropies are of similar order.

THEOREM 2.12 (THEOREM 3.2 [32]). Let $\rho_{XE} \in \text{Pos}(\mathcal{H}_X \otimes \mathcal{H}_E)$ be positive semidefinite of the form $\rho_{XE} = \sum_{x \in \mathcal{X}} |x\rangle\langle x| \otimes \rho_E^x$, where $\mathcal{X}$ is a finite alphabet. Let $\sigma_E \in \text{D}(\mathcal{H}_E)$ be an arbitrary density matrix. Then for any $\delta > 0$ and $0 < \varepsilon \leq 1$,

$$H_{\infty}^{\delta}(X|E)_{\rho} \geq -\frac{1}{\varepsilon} \log \left(\sum_x \tilde{Q}_{1+\varepsilon}(\rho_E^x \| \sigma_E) \right) - \frac{1 + 2 \log(1/\delta)}{\varepsilon}.$$

3 TRAPDOOR CLAW-FREE HASH FUNCTIONS

Let λ be a security parameter, and $\mathcal{X}$ and $\mathcal{Y}$ finite sets (depending on λ). For our purposes an ideal family of functions $\mathcal{F}$ would have the following properties: For each public key k , there are two functions $\{f_{k,b} : \mathcal{X} \rightarrow \mathcal{Y}\}_{b \in \{0,1\}}$ that are both injective and have the same range, and are invertible given a suitable trapdoor t_k (i.e., t_k can be used to compute x given b and $y = f_{k,b}(x)$). Furthermore, the pair of functions should be claw-free: It must be hard for an attacker to find two pre-images $x_0, x_1 \in \mathcal{X}$ such that $f_{k,0}(x_0) = f_{k,1}(x_1)$. Finally, the functions should satisfy an adaptive hardcore bit property, which is a stronger form of the claw-free property: Assuming for convenience that $\mathcal{X} = \{0, 1\}^w$, we would like that it is computationally infeasible to simultaneously generate a pair $(b, x_b) \in \{0, 1\} \times \mathcal{X}$ and a $d \in \{0, 1\}^w \setminus \{0^w\}$ such that with non-negligible advantage over $\frac{1}{2}$ the equation $d \cdot (x_0 \oplus x_1) = 0$, where x_{1-b} is defined as the unique element such that $f_{k,1-b}(x_{1-b}) = f_{k,b}(x_b)$, holds.

Unfortunately, we do not know how to construct a function family that exactly satisfies all these requirements under standard cryptographic assumptions. Instead, we construct a family that satisfies slightly relaxed requirements, which we will show still suffice for our purposes, based on the hardness of the learning with errors problem introduced in Section 2.3. The requirements are relaxed as follows: First, the range of the functions is no longer a set $\mathcal{Y}$; instead, it is $\mathcal{D}_{\mathcal{Y}}$, the set of probability densities over $\mathcal{Y}$. That is, each function returns a density, rather than a point. The trapdoor injective pair property is then described in terms of the support of the output densities: These supports should either be identical, for a colliding pair, or be disjoint, in all other cases.

The consideration of functions that return densities gives rise to an additional requirement of efficiency: There should exist a quantum polynomial-time procedure that efficiently prepares a superposition over the range of the function, i.e., for any key k and $b \in \{0, 1\}$, the procedure can prepare the state

$$\frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}} \sqrt{(f_{k,b}(x))(y)|x\rangle|y\rangle}. \quad (12)$$

In our instantiation based on LWE, it is not possible to prepare (12) perfectly, but it is possible to create a superposition with coefficients $\sqrt{(f'_{k,b}(x))(y)}$, such that the resulting state is within negligible trace distance of (12). The density $f'_{k,b}(x)$ is required to satisfy two properties used in our protocol. First, it must be easy to check, without the trapdoor, if an $y \in \mathcal{Y}$ lies in the support of $f'_{k,b}(x)$. Second, the inversion algorithm should operate correctly on all y in the support of $f'_{k,b}(x)$.

We slightly modify the adaptive hardcore bit requirement as well. Since the set $\mathcal{X}$ may not be a subset of binary strings, we first assume the existence of an injective, efficiently invertible map $J : \mathcal{X} \rightarrow \{0, 1\}^w$. Next, we only require the adaptive hardcore bit property to hold for a subset of all nonzero strings, instead of the set $\{0, 1\}^w \setminus \{0^w\}$. Finally, membership in the appropriate set should be efficiently checkable, given access to the trapdoor.

A formal definition follows.

Definition 3.1 (NTCF family). Let λ be a security parameter. Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets. Let $\mathcal{K}_{\mathcal{F}}$ be a finite set of keys. A family of functions

$$\mathcal{F} = \{f_{k,b} : \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}\}_{k \in \mathcal{K}_{\mathcal{F}}, b \in \{0,1\}}$$

is called a **noisy trapdoor claw free (NTCF) family** if the following conditions hold:

- (1) **Efficient Function Generation.** There exists an efficient probabilistic algorithm $\text{GEN}_{\mathcal{F}}$ that generates a key $k \in \mathcal{K}_{\mathcal{F}}$ together with a trapdoor t_k :

$$(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda).$$

- (2) **Trapdoor Injective Pair.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold:

- (a) *Trapdoor:* There exists an efficient deterministic algorithm $\text{INV}_{\mathcal{F}}$ such that for all $b \in \{0, 1\}$, $x \in \mathcal{X}$, and $y \in \text{SUPP}(f_{k,b}(x))$, $\text{INV}_{\mathcal{F}}(t_k, b, y) = x$. Note that this implies that for all $b \in \{0, 1\}$ and $x \neq x' \in \mathcal{X}$, $\text{SUPP}(f_{k,b}(x)) \cap \text{SUPP}(f_{k,b}(x')) = \emptyset$.

- (b) *Injective pair:* There exists a perfect matching $\mathcal{R}_k \subseteq \mathcal{X} \times \mathcal{X}$ such that $f_{k,0}(x_0) = f_{k,1}(x_1)$ if and only if $(x_0, x_1) \in \mathcal{R}_k$.

- (3) **Efficient Range Superposition.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ and $b \in \{0, 1\}$ there exists a function $f'_{k,b} : \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}$ such that the following hold:

- (a) For all $(x_0, x_1) \in \mathcal{R}_k$ and $y \in \text{SUPP}(f'_{k,b}(x_b))$, $\text{INV}_{\mathcal{F}}(t_k, b, y) = x_b$ and $\text{INV}_{\mathcal{F}}(t_k, b \oplus 1, y) = x_{b \oplus 1}$.

- (b) There exists an efficient deterministic procedure $\text{CHK}_{\mathcal{F}}$ that, on input $k, b \in \{0, 1\}$, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, returns 1 if $y \in \text{SUPP}(f'_{k,b}(x))$ and 0 otherwise. Note that $\text{CHK}_{\mathcal{F}}$ is not provided the trapdoor t_k .

- (c) For every k and $b \in \{0, 1\}$,

$$\mathbb{E}_{x \leftarrow \mathcal{X}} \left[H^2(f_{k,b}(x), f'_{k,b}(x)) \right] \leq \mu(\lambda),$$

for some negligible function $\mu(\cdot)$. Here, H^2 is the Hellinger distance; see (3). Moreover, there exists an efficient procedure $\text{SAMP}_{\mathcal{F}}$ that on input k and $b \in \{0, 1\}$ prepares the state

$$\frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{(f'_{k,b}(x))(y)|x\rangle|y\rangle}. \quad (13)$$

- (4) **Adaptive Hardcore Bit.** For all keys $k \in \mathcal{K}_{\mathcal{F}}$ the following conditions hold, for some integer w that is a polynomially bounded function of λ .

- (a) For all $b \in \{0, 1\}$ and $x \in \mathcal{X}$, there exists a set $G_{k,b,x} \subseteq \{0, 1\}^w$ such that $\Pr_{d \leftarrow \mathcal{U}_{\{0,1\}^w}}[d \notin G_{k,b,x}]$ is negligible, and moreover there exists an efficient algorithm that checks for membership in $G_{k,b,x}$ given k, b, x and the trapdoor t_k .

(b) There is an efficiently computable injection $J : \mathcal{X} \rightarrow \{0, 1\}^w$, such that J can be inverted efficiently on its range, and such that the following holds: If

$$H_k = \left\{ (b, x_b, d, d \cdot (J(x_0) \oplus J(x_1))) \mid b \in \{0, 1\}, (x_0, x_1) \in \mathcal{R}_k, d \in G_{k,0,x_0} \cap G_{k,1,x_1} \right\},^2$$

$$\overline{H}_k = \left\{ (b, x_b, d, c) \mid (b, x, d, c \oplus 1) \in H_k \right\},$$

then for any quantum polynomial-time procedure $\mathcal{A}$ there exists a negligible function $\mu(\cdot)$ such that

$$\left| \Pr_{(k,t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)} [\mathcal{A}(k) \in H_k] - \Pr_{(k,t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)} [\mathcal{A}(k) \in \overline{H}_k] \right| \leq \mu(\lambda). \quad (14)$$

4 A TRAPDOOR CLAW-FREE FAMILY BASED ON LWE

In this section, we present our LWE-based construction of an NTCF. For LWE-related preliminaries and definitions, see Section 2.3. Let λ be a security parameter. All other parameters are functions of λ . Let $q \geq 2$ be a prime. Let $\ell, n, m \geq 1$ be polynomially bounded functions of λ and B_L, B_V, B_P be positive integers such that the following conditions hold:

(A.1) $n = \Omega(\ell \log q + \lambda)$,

(A.2) $m = \Omega(n \log q)$,

(A.3) $B_P = \frac{q}{2^{C_T} \sqrt{mn \log q}}$, for C_T the universal constant in Theorem 2.6,

(A.4) We have $B_L < B_V < B_P$ so the ratios $\frac{B_P}{B_V}$ and $\frac{B_V}{B_L}$ are both super-polynomial in λ .

Given a choice of parameters satisfying all conditions (A.1) to (A.4), we describe the function family $\mathcal{F}_{\text{LWE}}$. Let $\mathcal{X} = \mathbb{Z}_q^n$ and $\mathcal{Y} = \mathbb{Z}_q^m$. The key space $\mathcal{K}_{\mathcal{F}_{\text{LWE}}}$ is a subset of $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ defined in Section 4.1. For $b \in \{0, 1\}$, $x \in \mathcal{X}$ and key $k = (A, As + e)$, the density $f_{k,b}(x)$ is defined as

$$\forall y \in \mathcal{Y}, \quad (f_{k,b}(x))(y) = D_{\mathbb{Z}_q^m, B_P}(y - Ax - b \cdot As), \quad (15)$$

where the density $D_{\mathbb{Z}_q^m, B_P}$ is defined in (6). It follows from the definition of the key generation procedure $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ given in Section 4.1 that $f_{k,b}$ is well-defined given $k = (A, As + e)$, as for our choice of parameters k uniquely identifies s .

The four properties required for a noisy trapdoor claw-free family, as specified in Definition 3.1, are verified in the following subsections, providing a proof of the following theorem. Recall the definition of the hardness assumption $\text{LWE}_{n,q,\chi}$ given in Definition 2.5.

THEOREM 4.1. *For any choice of parameters satisfying the conditions (A.1) to (A.4), the function family $\mathcal{F}_{\text{LWE}}$ is a noisy trapdoor claw free family under the hardness assumption $\text{LWE}_{\ell,q,D_{\mathbb{Z}_q^m, B_L}}$.*

Remark 4.2. We briefly discuss possible parameter settings for a correct and secure realization of the construction.

For known worst-case to average-case reductions to apply [37], we should set $B_L \geq 2\sqrt{\ell}$. For the sake of efficiency, we can choose B_L so equality holds. Since the evaluation algorithms run in $\text{poly}(\ell)$ time, we should take $\ell = \text{poly}(\lambda)$. The ratios $\frac{B_P}{B_V} = \frac{B_V}{B_L}$ affect the so-called “statistical security parameter” of the construction. Aiming for $2^{-\lambda}$ statistical security, we may set $\frac{B_P}{B_V} = \frac{B_V}{B_L} = 2^\lambda$.

Once ℓ has been chosen, the parameters n, m are determined by conditions (A.1), (A.2) and q is determined by condition (A.3). These conditions already imply that $q = 2^{2\lambda} / \text{poly}(\lambda)$. We need to set ℓ so the LWE problem with the resulting q is computationally hard. The hardness of the

²Note that although both x_0 and x_1 are referred to define the set H_k , only one of them, x_b , is explicitly specified in any 4-tuple that lies in H_k .

LWE problem scales very roughly as $2^{\tilde{\Omega}(\ell / \log(q/B_L))}$ (see, e.g., References [13, 39, 40]). In our case $\log(q/B_L) = O(\lambda)$ and, therefore, we can choose $\ell \approx \lambda^2$, which would imply exponential hardness (in λ).

We note that other choices of parameters are possible. For example, one could be satisfied with a statistical security parameter that is smaller than the computational security guarantee, thus choosing $\frac{B_P}{B_V}, \frac{B_V}{B_L}$ as more moderate functions of λ and improving efficiency. Another possible consideration is that in our suggested setting the ratio q/B_L scales sub-exponentially with ℓ , which corresponds to the hardness of sub-exponential approximation for lattice problems. One might not want to assume that sub-exponential approximation is hard and instead choose the parameters so q/B_L scales more moderately as a function of ℓ .

4.1 Efficient Function Generation

$\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ is defined as follows: First, the procedure samples a random $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, together with trapdoor information $t_{\mathbf{A}}$. This is done using the procedure $\text{GENTRAP}(1^n, 1^m, q)$ from Theorem 2.6. Recall that Assumption (A.2) requires that $m = \Omega(n \log q)$ as needed for the theorem to hold. The trapdoor allows the evaluation of an inversion algorithm INVERT that, on input $\mathbf{A}$, $t_{\mathbf{A}}$ and $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ returns $\mathbf{s}$ and $\mathbf{e}$ as long as $\|\mathbf{e}\| \leq \frac{q}{C_T \sqrt{n \log q}}$. Moreover, the distribution on matrices $\mathbf{A}$ returned by GENTRAP is negligibly close to the uniform distribution on $\mathbb{Z}_q^{m \times n}$.

Next, the sampling procedure selects $\mathbf{s} \in \{0, 1\}^n$ uniformly at random, and a vector $\mathbf{e} \in \mathbb{Z}_q^m$ by sampling each coordinate independently according to the distribution $D_{\mathbb{Z}_q, B_V}$ defined in (6). $\text{GEN}_{\mathcal{F}_{\text{LWE}}}$ returns $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and $t_k = t_{\mathbf{A}}$.

4.2 Trapdoor Injective Pair

- (a) *Trapdoor.* It follows from (15) and the definition of the distribution $D_{\mathbb{Z}_q^m, B_P}$ in (6) that for any key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$ and for all $x \in \mathcal{X}$,

$$\text{SUPP}(f_{k,0}(x)) = \{\mathbf{A}x + \mathbf{e}_0 \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m}\}, \quad (16)$$

$$\text{SUPP}(f_{k,1}(x)) = \{\mathbf{A}x + \mathbf{A}\mathbf{s} + \mathbf{e}_0 \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m}\}. \quad (17)$$

The procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ takes as input the trapdoor $t_{\mathbf{A}}$, $b \in \{0, 1\}$, and $y \in \mathcal{Y}$. It uses the algorithm INVERT to determine $\mathbf{s}_0, \mathbf{e}_0$ such that $y = \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$, and returns the element $\mathbf{s}_0 - b \cdot \mathbf{s} \in \mathcal{X}$. Using Theorem 2.6, this procedure returns the unique correct outcome provided $y = \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$ for some $\mathbf{e}_0$ such that $\|\mathbf{e}_0\| \leq \frac{q}{C_T \sqrt{n \log q}}$. This condition is satisfied for all $y \in \text{SUPP}(f_{k,b}(x))$ provided B_P is chosen so

$$B_P \leq \frac{q}{C_T \sqrt{mn \log q}}, \quad (18)$$

which is satisfied by the choice in (A.3).

- (b) *Injective Pair.* We let $\mathcal{R}_k$ be the set of all pairs (x_0, x_1) such that $f_{k,0}(x_0) = f_{k,1}(x_1)$. By definition, this occurs if and only if $x_1 = x_0 - \mathbf{s}$, and so $\mathcal{R}_k$ is a perfect matching.

4.3 Efficient Range Superposition

For $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \in \mathcal{K}_{\mathcal{F}_{\text{LWE}}}$, $b \in \{0, 1\}$ and $x \in \mathcal{X}$, let

$$(f'_{k,b}(x))(y) = D_{\mathbb{Z}_q^m, B_P}(y - \mathbf{A}x - b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e})). \quad (19)$$

Note that $f'_{k,0}(x) = f_{k,0}(x)$ for all $x \in \mathcal{X}$. The distributions $f'_{k,1}(x)$ and $f_{k,1}(x)$ are shifted by $\mathbf{e}$. Given the key k and $x \in \mathcal{X}$, the densities $f'_{k,0}(x)$ and $f'_{k,1}(x)$ are efficiently computable. For all

$x \in \mathcal{X}$,

$$\text{SUPP}(f'_{k,0}(x)) = \text{SUPP}(f_{k,0}(x)), \quad (20)$$

$$\text{SUPP}(f'_{k,1}(x)) = \{ \mathbf{A}x + \mathbf{e}_0 + \mathbf{A}\mathbf{s} + \mathbf{e} \mid \|\mathbf{e}_0\| \leq B_P \sqrt{m} \}. \quad (21)$$

- (a) Using that $B_V < B_P$, it follows that the norm of the term $\mathbf{e}_0 + \mathbf{e}$ in (21) is always at most $2B_P \sqrt{m}$. Therefore, the inversion procedure $\text{INV}_{\mathcal{F}_{\text{LWE}}}$ can be guaranteed to return x on input $t_A, b \in \{0, 1\}, y \in \text{SUPP}(f'_{k,b}(x))$ if we strengthen the requirement on B_P given in (18) to

$$B_P \leq \frac{q}{2C_T \sqrt{mn \log q}}, \quad (22)$$

which is still satisfied by (A.3). This strengthened trapdoor requirement also implies that for all $b \in \{0, 1\}, (x_0, x_1) \in \mathcal{R}_k$, and $y \in \text{SUPP}(f'_{k,b}(x_b))$, $\text{INV}_{\mathcal{F}_{\text{LWE}}}(t_A, b \oplus 1, y) = x_{b \oplus 1}$.

- (b) On input $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$, $b \in \{0, 1\}$, $x \in \mathcal{X}$, and $y \in \mathcal{Y}$, the procedure $\text{CHK}_{\mathcal{F}_{\text{LWE}}}$ operates as follows: If $b = 0$, then it computes $\mathbf{e}' = y - \mathbf{A}x$. If $\|\mathbf{e}'\| \leq B_P \sqrt{m}$, then the procedure returns 1, and 0 otherwise. If $b = 1$, then it computes $\mathbf{e}' = y - \mathbf{A}x - (\mathbf{A}\mathbf{s} + \mathbf{e})$. If $\|\mathbf{e}'\| \leq B_P \sqrt{m}$, then it returns 1, and 0 otherwise.
- (c) We bound the Hellinger distance between the densities $f_{k,b}(x)$ and $f'_{k,b}(x)$. If $b = 0$, then they are identical. If $b = 1$, then both densities are shifts of $D_{\mathbb{Z}_q^m, B_P}$, where the shifts differ by $\mathbf{e}$. Each coordinate of $\mathbf{e}$ is drawn independently from $D_{\mathbb{Z}_q, B_V}$, so $\|\mathbf{e}\| \leq \sqrt{m}B_V$. Applying Lemma 2.4, we get that

$$H^2(f_{k,1}(x), f'_{k,1}(x)) \leq 1 - e^{-\frac{2\pi m B_V}{B_P}}.$$

Using the assumption that B_P/B_V is super-polynomial as required in Assumption (A.4), this is negligible, as desired. It remains to describe the procedure $\text{SAMP}_{\mathcal{F}_{\text{LWE}}}$. At the first step, the procedure creates the following superposition:

$$\sum_{\mathbf{e}_0 \in \mathbb{Z}_q^m} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |\mathbf{e}_0\rangle. \quad (23)$$

This state can be prepared efficiently as described in Reference [37, Lemma 3.12].³

At the second step, the procedure creates a uniform superposition over $x \in \mathcal{X}$, yielding the state

$$q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |x\rangle |\mathbf{e}_0\rangle. \quad (24)$$

At the third step, using the key $k = (\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ and the input bit b the procedure computes

$$q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0)} |x\rangle |\mathbf{e}_0\rangle |\mathbf{A}x + \mathbf{e}_0 + b \cdot (\mathbf{A}\mathbf{s} + \mathbf{e})\rangle. \quad (25)$$

³Specifically, the state can be created using a technique by Grover and Rudolph ([23]), who show that to create such a state, it suffices to have the ability to efficiently compute the sum $\sum_{x=c}^d D_{\mathbb{Z}_q, B_P}(x)$ for any $c, d \in \{-\lfloor \sqrt{B_P} \rfloor, \dots, \lceil \sqrt{B_P} \rceil\} \subseteq \mathbb{Z}_q$ and to within good precision. This can be done using standard techniques used to sample from the normal distribution.

At this point, observe that $\mathbf{e}_0$ can be computed from x , the last register, b and the key k . The procedure can then uncompute the register containing $\mathbf{e}_0$, yielding

$$\begin{aligned}
 & q^{-\frac{n}{2}} \sum_{\substack{x \in \mathcal{X} \\ \mathbf{e}_0 \in \mathbb{Z}_q^m}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(\mathbf{e}_0) |x\rangle |A\mathbf{x} + \mathbf{e}_0 + b \cdot (A\mathbf{s} + \mathbf{e})\rangle} \\
 &= q^{-\frac{n}{2}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{D_{\mathbb{Z}_q^m, B_P}(y - A\mathbf{x} - b \cdot (A\mathbf{s} + \mathbf{e})) |x\rangle |y\rangle} \\
 &= q^{-\frac{n}{2}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \sqrt{(f'_{k,b}(x))(y) |x\rangle |y\rangle}. \tag{26}
 \end{aligned}$$

4.4 Adaptive Hardcore Bit

This section is devoted to the proof that condition 4 of Definition 3.1 holds. We start by providing a formal statement. Recall that $\mathcal{X} = \mathbb{Z}_q^n$ and let $w = n \lceil \log q \rceil$. Let $J : \mathcal{X} \rightarrow \{0, 1\}^w$ be such that $J(x)$ returns the binary representation of $x \in \mathcal{X}$. For $b \in \{0, 1\}$, $x \in \mathcal{X}$, and $d \in \{0, 1\}^w$, let $I_{b,x}(d) \in \{0, 1\}^n$ be the vector whose each coordinate is obtained by taking the inner product mod 2 of the corresponding block of $\lceil \log q \rceil$ coordinates of d and of $J(x) \oplus J(x - (-1)^b \mathbf{1})$, where $\mathbf{1} \in \mathbb{Z}_q^n$ is the vector with all its coordinates equal to 1 $\in \mathbb{Z}_q$. For $k = (A, A\mathbf{s} + \mathbf{e})$, $b \in \{0, 1\}$ and $x \in \mathcal{X}$, we define the set $G_{k,b,x}$ as

$$G_{k,b,x} = \left\{ d \in \{0, 1\}^w \mid \exists i \in \left\{ b \frac{n}{2}, \dots, b \frac{n}{2} + \frac{n}{2} \right\} : (I_{b,x}(d))_i \neq 0 \right\}.$$

Observe that for all $b \in \{0, 1\}$ and $x \in \mathcal{X}$, if d is sampled uniformly at random, $d \notin G_{k,b,x}$ with negligible probability. This follows simply because for any $b \in \{0, 1\}$, $J(x) \oplus J(x - (-1)^b \mathbf{1})$ is non-zero, since J is injective. Observe also that checking membership in $G_{k,b,x}$ is possible given only b, x . This shows condition 4.(a) in the adaptive hardcore bit condition in Definition 3.1.

Given $(x_0, x_1) \in \mathcal{R}_k$ (where $k = (A, A\mathbf{s} + \mathbf{e})$), recall from Section 4.2 that $x_1 = x_0 - \mathbf{s}$. We use the following notation: We write $s \in \{0, 1\}^n$ as $s = (s_0, s_1)$, where $s_0, s_1 \in \{0, 1\}^{\frac{n}{2}}$ are the $\frac{n}{2}$ -bit prefix and suffix of s , respectively (for simplicity, assume n is even; if not, then ties can be broken arbitrarily). For convenience, we also introduce the following set, where $y = f_{k,0}(x_0) = f_{k,1}(x_1)$:

$$\hat{G}_{s_1, 0, x_0} = \hat{G}_{s_0, 1, x_1} = G_{k, 0, x_0} \cap G_{k, 1, x_1}. \tag{27}$$

The motivation for using two different notations for the same set is to clarify that membership in the set can be decided given $(s_{b \oplus 1}, b, x_b)$, for either $b \in \{0, 1\}$. This point will be important in the proof of Lemma 4.4.

The following lemma establishes item 4.(b) in Definition 3.1:

LEMMA 4.3. Assume a choice of parameters satisfying the conditions (A.1) to (A.4). Assume the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ holds. Let $s \in \{0, 1\}^n$. Let⁴

$$H_s = \left\{ (b, x, d, d \cdot (J(x) \oplus J(x - (-1)^b \mathbf{s}))) \mid b \in \{0, 1\}, x \in \mathcal{X}, d \in \hat{G}_{s_{b \oplus 1}, b, x} \right\}, \tag{28}$$

$$\overline{H}_s = \left\{ (b, x, d, c) \mid (b, x, d, c \oplus 1) \in H_s \right\}. \tag{29}$$

Then for any quantum polynomial-time procedure

$$\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0, 1\} \times \mathcal{X} \times \{0, 1\}^w \times \{0, 1\}$$

⁴We write the sets as H_s instead of H_k to emphasize the dependence on s .

there exists a negligible function $\mu(\lambda)$ such that

$$\left| \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As+e) \in H_s] - \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As+e) \in \overline{H_s}] \right| \leq \mu(\lambda). \quad (30)$$

The proof of the lemma proceeds in three steps. First, in Section 4.4.2, we establish some preliminary results on the distribution of the inner product $(\hat{d} \cdot s \bmod 2)$, where $\hat{d} \in \{0, 1\}^n$ is a fixed nonzero binary vector and $s \leftarrow_U \{0, 1\}^n$ a uniformly random binary vector, conditioned on $Cs = v$ for some randomly chosen matrix $C \in \mathbb{Z}_q^{\ell \times n}$ and arbitrary $v \in \mathbb{Z}_q^\ell$. This condition is combined with the LWE assumption in Section 4.4.3 to argue that $(\hat{d} \cdot s \bmod 2)$ remains computationally indistinguishable from uniform even when the matrix C is an LWE matrix A , and the adversary is able to choose $\hat{d}$ after being given access to $As + e$ for some error vector $e \in \mathbb{Z}_q^m$. This will allow us to derive the following lemma, whose proof is provided in Section 4.4.3.

We will show computational indistinguishability based on the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ specified in Definition 2.5. Since our goal is to prove Lemma 4.3, we consider procedures that output a tuple $(b, x, d, c) \in \{0, 1\} \times \mathcal{X} \times \{0, 1\}^w \times \{0, 1\}$.

LEMMA 4.4. Assume a choice of parameters satisfying the conditions (A.1) to (A.4). Assume the hardness assumption $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ holds. Let

$$\mathcal{A} : \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m \rightarrow \{0, 1\} \times \mathcal{X} \times \{0, 1\}^w \times \{0, 1\}$$

be a quantum polynomial-time procedure. For $b \in \{0, 1\}$ and $x \in \mathcal{X}$, let $I_{b,x} : \{0, 1\}^w \rightarrow \{0, 1\}^n$ be an efficiently computable map. For every $s = (s_0, s_1) \in \{0, 1\}^n$ and $(b, x) \in \{0, 1\} \times \mathcal{X}$, let $\hat{G}_{s_{b \oplus 1}, b, x} \subseteq \{0, 1\}^w$ be a set depending only on b, x and $s_{b \oplus 1}$ and such that for all $d \in \hat{G}_{s_{b \oplus 1}, b, x}$ the first (if $b = 0$) or last (if $b = 1$) $\frac{n}{2}$ bits of $I_{b,x}(d)$ are not all 0. Then the distributions

$$D_0 = \left((A, As + e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda), (b, x, d, c) \leftarrow \mathcal{A}(A, As + e), I_{b,x}(d) \cdot s \bmod 2 \right) \quad (31)$$

and

$$D_1 = \left((A, As + e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda), (b, x, d, c) \leftarrow \mathcal{A}(A, As + e), (\delta_{d \in \hat{G}_{s_{b \oplus 1}, b, x}} r) \oplus (I_{b,x}(d) \cdot s \bmod 2) \right), \quad (32)$$

where $r \leftarrow_U \{0, 1\}$ and $\delta_{d \in \hat{G}_{s_{b \oplus 1}, b, x}}$ is 1 if $d \in \hat{G}_{s_{b \oplus 1}, b, x}$ and 0 otherwise, are computationally indistinguishable.

We prove Lemma 4.3 from Lemma 4.4 by relating the inner product appearing in the definition of H_k (in condition 4.(b) of Definition 3.1) to an inner product of the form $\hat{d} \cdot s$, where $\hat{d}$ can be efficiently computed from d . This proof appears in Section 4.4.1 below.

4.4.1 Proof of Lemma 4.3 from Lemma 4.4. The proof is by contradiction. Assume that there exists a quantum polynomial-time procedure $\mathcal{A}$ such that the left-hand side of (30) is at least some non-negligible function $\eta(\lambda)$. We derive a contradiction by showing that the two distributions D_0 and D_1 in Lemma 4.4, for $I_{b,x}$ defined at the start of this section and $\hat{G}_{s_{b \oplus 1}, b, x}$ defined in (27), are computationally distinguishable, giving a contradiction.

Let $(A, As + e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)$ and $(b, x, d, c) \leftarrow \mathcal{A}(A, As + e)$. To link $\mathcal{A}$ to the distributions in Lemma 4.4, we relate the inner product condition in (28) to an inner product $\hat{d} \cdot s$ of the form appearing in (31), for $\hat{d} = I_{b,x}(d)$ that can be computed efficiently from b, x , and d . This is based on the following claim:

CLAIM 4.5. For all $b \in \{0, 1\}, x \in X, d \in \{0, 1\}^w$ and $s \in \{0, 1\}^n$ the following equality holds:

$$d \cdot (J(x) \oplus J(x - (-1)^b s)) = I_{b,x}(d) \cdot s. \quad (33)$$

Moreover, the function $I_{b,x}$ is efficiently computable given b, x .

PROOF. We do the proof in case $n = 1$ and $w = \lceil \log q \rceil$, as the case of general n follows by linearity. In this case s is a single bit. If $s = 0$, then both sides of (33) evaluate to zero, so the equality holds trivially. It then suffices to define $I_{b,x_b}(d)$ so the equation holds when $s = 1$. A choice of either of

$$I_{0,x_0}(d) = d \cdot (J(x_0) \oplus J(x_0 - 1)), \quad I_{1,x_1}(d) = d \cdot (J(x_1) \oplus J(x_1 + 1))$$

satisfies all requirements. It is clear from the definition of $I_{b,x}$ that it can be computed efficiently given b, x . $\square$

The procedure $\mathcal{A}$, the function $I_{b,x}$ defined at the start of this section, and the sets $\hat{G}_{s_{b \oplus 1}, b, x}$ in (27) fully specify D_0 and D_1 . To conclude, we construct a distinguisher $\mathcal{A}'$ between D_0 and D_1 . Consider two possible distinguishers, $\mathcal{A}'_u$ for $u \in \{0, 1\}$. Given a sample $w = ((A, As + e), (b, x, d, c), t)$, $\mathcal{A}'_u$ returns 0 if $c = t \oplus u$, and 1 otherwise. First note that

$$\begin{aligned} & \sum_{u \in \{0, 1\}} \left| \Pr_{w \leftarrow D_0} [\mathcal{A}'_u(w) = 0] - \Pr_{w \leftarrow D_1} [\mathcal{A}'_u(w) = 0] \right| \\ &= \sum_{u \in \{0, 1\}} \left| \Pr_{w \leftarrow D_0} [\mathcal{A}'_u(w) = 0 \wedge d \in \hat{G}_{s_{b \oplus 1}, b, x}] - \Pr_{w \leftarrow D_1} [\mathcal{A}'_u(w) = 0 \wedge d \in \hat{G}_{s_{b \oplus 1}, b, x}] \right|, \end{aligned} \quad (34)$$

since if $d \notin \hat{G}_{s_{b \oplus 1}, b, x}$, then the distributions D_0 and D_1 are identical by definition. Next, if the sample held by $\mathcal{A}'_u$ is from the distribution D_0 and if $(b, x, d, c) \in H_s$, then by the definition of H_s and (33) it follows that $c = d \cdot (J(x) \oplus J(x - (-1)^b s)) = I_{b,x}(d) \cdot s = t$. If instead $(b, x, d, c) \in \bar{H}_s$, then $c \oplus 1 = d \cdot (J(x) \oplus J(x - (-1)^b s)) = I_{b,x}(d) \cdot s = t$. The expression in (34) is thus equal to:

$$\begin{aligned} (34) &= \left| \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As + e) \in H_s] - \frac{1}{2} \Pr_{w \leftarrow D_1} [d \in \hat{G}_{s_{b \oplus 1}, b, x}] \right| \\ &\quad + \left| \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As + e) \in \bar{H}_s] - \frac{1}{2} \Pr_{w \leftarrow D_1} [d \in \hat{G}_{s_{b \oplus 1}, b, x}] \right| \\ &\geq \left| \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As + e) \in H_s] - \Pr_{(A, As+e) \leftarrow \text{GEN}_{\mathcal{F}_{\text{LWE}}}(1^\lambda)} [\mathcal{A}(A, As+e) \in \bar{H}_s] \right| \\ &\geq \eta. \end{aligned}$$

Therefore, at least one of $\mathcal{A}'_0$ or $\mathcal{A}'_1$ must successfully distinguish between D_0 and D_1 with advantage at least $\frac{\eta}{2}$, a contradiction with the statement of Lemma 4.4. $\square$

4.4.2 A Building Block: Moderate Matrices. The following lemma argues that, provided the matrix $C \in \mathbb{Z}_q^{\ell \times n}$ is a uniformly random matrix with sufficiently few rows, the distribution (C, Cs) for arbitrary $s \in \{0, 1\}^n$ does not reveal any parity of s .

LEMMA 4.6. Let q be a prime, $\ell, n \geq 1$ integers, and $C \in \mathbb{Z}_q^{\ell \times n}$ a uniformly random matrix. With probability at least $1 - q^{-\ell} \cdot 2^{-\frac{n}{8}}$ over the choice of C the following holds: For a fixed C , all $v \in \mathbb{Z}_q^\ell$ and $\hat{d} \in \{0, 1\}^n \setminus \{0^n\}$, the distribution of $(\hat{d} \cdot s \bmod 2)$, where s is uniform in $\{0, 1\}^n$ conditioned on $Cs = v$, is within statistical distance $O(q^{\frac{3\ell}{2}} \cdot 2^{-\frac{n}{40}})$ of the uniform distribution over $\{0, 1\}$.

To prove the lemma, we introduce the notion of a *moderate* matrix.

Definition 4.7. Let $\mathbf{b} \in \mathbb{Z}_q^n$. We say that $\mathbf{b}$ is moderate if it contains at least $\frac{n}{4}$ entries whose unique representative in $(-q/2, q/2]$ has its absolute value in the range $(\frac{q}{8}, \frac{3q}{8}]$. A matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ is moderate if its entire row span (except 0^n) is moderate.

LEMMA 4.8. Let q be prime and ℓ, n be integers. Then

$$\Pr_{\mathbf{C} \leftarrow \mathcal{U}(\mathbb{Z}_q^{\ell \times n})} (\mathbf{C} \text{ is moderate}) \geq 1 - q^\ell \cdot 2^{-\frac{n}{8}}.$$

PROOF. Consider an arbitrary non-zero vector $\mathbf{b}$ in the row-span of a uniform $\mathbf{C}$. Then the marginal distribution of $\mathbf{b}$ is uniform. By Chernoff, $\mathbf{b}$ is moderate with probability at least $1 - e^{-\frac{2n}{16}} \geq 1 - 2^{-\frac{n}{8}}$. Applying the union bound over all at most $q^\ell - 1$ non-zero vectors in the row span, the result follows. $\square$

LEMMA 4.9. Let $\mathbf{C} \in \mathbb{Z}_q^{\ell \times n}$ be an arbitrary moderate matrix and let $\hat{\mathbf{d}} \in \{0, 1\}^n \setminus \{0^n\}$ be an arbitrary non-zero binary vector. Let $\mathbf{s}$ be uniform over $\{0, 1\}^n$ and consider the random variables $\mathbf{v} = \mathbf{C}\mathbf{s} \bmod q$ and $z = \hat{\mathbf{d}} \cdot \mathbf{s} \bmod 2$. Then $(\mathbf{v}, z)$ is within total variation distance at most $q^{\frac{\ell}{2}} \cdot 2^{-\frac{n}{40}}$ of the uniform distribution over $\mathbb{Z}_q^\ell \times \{0, 1\}$.

PROOF. Let f be the probability density function of $(\mathbf{v}, z)$. Interpreting z as an element of $\mathbb{Z}_2$, let $\hat{f}$ be the Fourier transform over $\mathbb{Z}_q^\ell \times \mathbb{Z}_2$. Let U denote the density of the uniform distribution over $\mathbb{Z}_q^\ell \times \mathbb{Z}_2$. Applying the Cauchy-Schwarz inequality,

$$\begin{aligned} \frac{1}{2} \|f - U\|_1 &\leq \sqrt{\frac{q^\ell}{2}} \|f - U\|_2 \\ &= \frac{1}{2} \|\hat{f} - \hat{U}\|_2 \\ &= \frac{1}{2} \left(\sum_{(\hat{\mathbf{v}}, \hat{z}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_2 \setminus \{(0, 0)\}} |\hat{f}(\hat{\mathbf{v}}, \hat{z})|^2 \right)^{1/2}, \end{aligned} \quad (35)$$

where the second line follows from Parseval's identity, and for the third line, we used $\hat{f}(0, 0) = \hat{U}(0, 0) = 1$ and $\hat{U}(\hat{\mathbf{v}}, \hat{z}) = 0$ for all $(\hat{\mathbf{v}}, \hat{z}) \neq (0^\ell, 0)$. To bound (35), we estimate the Fourier coefficients of f . Denoting $\omega_{2q} = e^{-\frac{2\pi i}{2q}}$, for any $(\hat{\mathbf{v}}, \hat{z}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_2$, we can write

$$\begin{aligned} \hat{f}(\hat{\mathbf{v}}, \hat{z}) &= \mathbb{E}_{\mathbf{s}} \left[\omega_{2q}^{(2 \cdot \hat{\mathbf{v}}^T \mathbf{C} + q \cdot \hat{z} \hat{\mathbf{d}}^T) \mathbf{s}} \right] \\ &= \mathbb{E}_{\mathbf{s}} \left[\omega_{2q}^{\mathbf{w}^T \mathbf{s}} \right] \\ &= \prod_i \mathbb{E}_{s_i} \left[\omega_{2q}^{w_i s_i} \right], \end{aligned} \quad (36)$$

where we wrote $\mathbf{w}^T = 2 \cdot \hat{\mathbf{v}}^T \mathbf{C} + q \cdot \hat{z} \hat{\mathbf{d}}^T \in \mathbb{Z}_{2q}^n$. It follows that $\hat{f}(0^\ell, 1) = 0$, since $(\mathbf{d} \cdot \mathbf{s} \bmod 2)$ is uniform for $\mathbf{s}$ uniform.

We now observe that for all $i \in \{1, \dots, n\}$ such that the representative of $(\hat{\mathbf{v}}^T \mathbf{C})_i$ in $(-q/2, q/2]$ has its absolute value in $(\frac{q}{8}, \frac{3q}{8}]$ it holds that $\frac{w_i}{q} \in (\frac{1}{4}, \frac{3}{4}] \bmod 1$, in which case

$$\left| \mathbb{E}_{s_i} [\omega_{2q}^{w_i s_i}] \right| = \left| \cos \left(\frac{\pi}{2} \cdot \frac{w_i}{q} \right) \right| \leq \cos \left(\frac{\pi}{8} \right) \leq 2^{-\frac{1}{10}}. \quad (37)$$

Since $\mathbf{C}$ is moderate, there are at least $\frac{n}{4}$ such entries, so from (36) it follows that $|\hat{f}(\hat{\mathbf{v}}, \hat{z})| \leq 2^{-\frac{n}{40}}$ for all $\hat{\mathbf{v}} \neq 0$. Recalling (35), the lemma is proved. $\square$

We now prove Lemma 4.6 by generalizing Lemma 4.9 to adaptive d (i.e., d can depend on C, C_s).

PROOF OF LEMMA 4.6. We assume C is moderate; by Lemma 4.8, C is moderate with probability at least $1 - q^\ell \cdot 2^{-\frac{n}{8}}$. Let s be uniform over $\{0, 1\}^n$, $D_1 = (Cs, \hat{d} \cdot s \bmod 2)$, and D_2 uniformly distributed over $\mathbb{Z}_q^\ell \times \{0, 1\}$. Using that C is moderate, it follows from Lemma 4.9 that

$$\varepsilon = \|D_1 - D_2\|_{TV} \leq q^{\frac{\ell}{2}} \cdot 2^{-\frac{n}{40}}. \quad (38)$$

Fix $\mathbf{v}_0 \in \mathbb{Z}_q^\ell$ and let

$$\Delta = \frac{1}{2} \sum_{b \in \{0, 1\}} \left| \Pr_{s \leftarrow_U \{0, 1\}^n} (\hat{d} \cdot s \bmod 2 = b \mid C\mathbf{s} = \mathbf{v}_0) - \frac{1}{2} \right|. \quad (39)$$

To prove the lemma, it suffices to establish the appropriate upper bound on Δ , for all $\mathbf{v}_0$. By definition,

$$\begin{aligned} \varepsilon = \|D_1 - D_2\|_{TV} &= \frac{1}{2} \sum_{b \in \{0, 1\}, \mathbf{v} \in \mathbb{Z}_q^\ell} \left| \Pr(C\mathbf{s} = \mathbf{v}) \Pr(\hat{d} \cdot s \bmod 2 = b \mid C\mathbf{s} = \mathbf{v}) - \frac{1}{2q^\ell} \right| \\ &\geq \frac{1}{2} \sum_{b \in \{0, 1\}} \left| \Pr(C\mathbf{s} = \mathbf{v}_0) \Pr(\hat{d} \cdot s \bmod 2 = b \mid C\mathbf{s} = \mathbf{v}_0) - \frac{1}{2q^\ell} \right| \\ &= \frac{1}{2} \sum_{b \in \{0, 1\}} \left| \Pr(C\mathbf{s} = \mathbf{v}_0) \left(\frac{1}{2} + (-1)^b \Delta \right) - \frac{1}{2q^\ell} \right|, \end{aligned} \quad (40)$$

where all probabilities are under a uniform choice of $s \leftarrow_U \{0, 1\}^n$, and the last line follows from the definition of Δ in (39). Applying the inequality $|a| + |b| \geq \max(|a - b|, |a + b|)$, valid for any real a, b , to (40) it follows that

$$\Pr(C\mathbf{s} = \mathbf{v}_0) \cdot \Delta \leq \varepsilon \quad \text{and} \quad \Pr(C\mathbf{s} = \mathbf{v}_0) \geq \frac{1}{q^\ell} - 2\varepsilon. \quad (41)$$

If $q^{3\ell/2} 2^{-\frac{n}{40}} > \frac{1}{3}$, then the bound claimed in the lemma is trivial. If $q^{3\ell/2} 2^{-\frac{n}{40}} \leq \frac{1}{3}$, then $\varepsilon q^\ell \leq \frac{1}{3}$, so it follows from (41) that $\Delta \leq 3q^\ell \varepsilon$, which together with (38) proves the lemma. $\square$

4.4.3 Proof of Lemma 4.4. We use Lemma 4.6 and prove computational indistinguishability by introducing a sequence of hybrids. For the first step, we let

$$D^{(1)} = \left((\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), I_{b,x}(d) \cdot s \bmod 2 \right), \quad (42)$$

where $\tilde{\mathbf{A}} = \mathbf{BC} + \mathbf{F} \leftarrow \text{LOSSY}(1^n, 1^m, 1^\ell, q, D_{\mathbb{Z}_q, B_L})$ is sampled from a lossy sampler (see Definition 2.7). From the definition, $\mathbf{F} \in \mathbb{Z}_q^{m \times n}$ has entries i.i.d. from the distribution $D_{\mathbb{Z}_q, B_L}$ over $\mathbb{Z}_q$. To see that D_0 and $D^{(1)}$ are computationally indistinguishable, first note that the distribution of matrices $\mathbf{A}$ generated by $\text{GEN}_{\mathcal{T}_{\text{LWE}}}$ is negligibly far from the uniform distribution (see Theorem 2.6). Next, by Theorem 2.8, under the $\text{LWE}_{\ell, q, D_{\mathbb{Z}_q, B_L}}$ assumption a uniformly random matrix $\mathbf{A}$ and a lossy matrix $\tilde{\mathbf{A}}$ are computationally indistinguishable. Note that this step, as well as subsequent steps, uses that $\mathcal{A}$ and $I_{b,x}$ are efficiently computable.

For the second step, we remove the term $\mathbf{Fs}$ from the lossy LWE sample $\tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}$ to obtain the distribution

$$D^{(2)} = \left((\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{BC} + \mathbf{F}, \mathbf{BC}\mathbf{s} + \mathbf{e}), I_{b,x}(d) \cdot s \bmod 2 \right). \quad (43)$$

Using that $\mathbf{s}$ is binary and the entries of $\mathbf{F}$ are taken from a B_L -bounded distribution, it follows that $\|\mathbf{F}\mathbf{s}\| \leq n\sqrt{m}B_L$. Applying Lemma 2.4, the statistical distance between $D^{(1)}$ and $D^{(2)}$ is at most

$$\gamma = \sqrt{2} \left(1 - e^{\frac{-2\pi mnB_L}{B_V}} \right)^{1/2}, \quad (44)$$

which is negligible, due to the requirement that $\frac{B_V}{B_L}$ is superpolynomial given in (A.4).

For the third step, observe that the distribution $D^{(2)}$ in (43) only depends on s_b through $\mathbf{C}$ s and $I_{b,x}(d) \cdot s$, where $\mathbf{C}$ is uniformly random. It follows from Lemma 4.6 that provided $\frac{n}{2} = \Omega(\ell \log q + \lambda)$ as required by Assumption (A.1), with overwhelming probability (in the security parameter) over the choice of $\mathbf{C}$, if we fix all variables except for s_b , the distribution of $(I_{b,x}(d) \cdot s \bmod 2)$ is statistically indistinguishable (within statistical distance $2^{-\lambda}$) from $r \leftarrow_U \{0, 1\}$ as long as the $\frac{n}{2}$ bits of $I_{b,x}(d)$ associated with s_b are not all 0 (i.e., the first $\frac{n}{2}$ bits if $b = 0$ or the last $\frac{n}{2}$ bits if $b = 1$). Using that for $d \in \hat{G}_{s_{b \oplus 1}, b, x}$ the $\frac{n}{2}$ bits of $I_{b,x}(d)$ associated with s_b are not all 0, the distribution $D^{(2)}$ in (43) is statistically indistinguishable from

$$D^{(3)} = \left((\mathbf{BC} + \mathbf{F}, \mathbf{BCs} + \mathbf{e}), (b, x, d, c) \leftarrow \mathcal{A}(\mathbf{BC} + \mathbf{F}, \mathbf{BCs} + \mathbf{e}), (\delta_{d \in \hat{G}_{s_{b \oplus 1}, b, x}} r) \oplus (I_{b,x}(d) \cdot s \bmod 2) \right),$$

where $r \leftarrow_U \{0, 1\}$.

For the fourth step, we reinsert the term $\mathbf{F}\mathbf{s}$ to obtain

$$D^{(4)} = \left(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}, (b, x, d, c) \leftarrow \mathcal{A}(\tilde{\mathbf{A}}, \tilde{\mathbf{A}}\mathbf{s} + \mathbf{e}), (\delta_{d \in \hat{G}_{s_{b \oplus 1}, b, x}} r) \oplus (I_{b,x}(d) \cdot s \bmod 2) \right).$$

Statistical indistinguishability between $D^{(3)}$ and $D^{(4)}$ follows similarly as between $D^{(1)}$ and $D^{(2)}$. Finally, computational indistinguishability between $D^{(4)}$ and D_1 follows similarly to between $D^{(1)}$ and D_0 . $\square$

5 PROTOCOL DESCRIPTION

We introduce two protocols. The first we call the (*general*) *randomness expansion protocol*, or Protocol 1. This is our main randomness expansion protocol. It is introduced in Section 5.1 and summarized in Figure 2. The protocol describes the interaction between a *verifier* and *prover*. Ultimately, we aim to obtain the guarantee that any computationally bounded prover that is accepted with non-negligible probability by the verifier in the protocol must generate transcripts that contain information-theoretic randomness.

The second protocol is called the *simplified protocol*, or Protocol 2. It is introduced in Section 5.2, and summarized in Figure 3. This protocol abstracts some of the main features Protocol 1 and will be used as a tool in the analysis (it is not meant to be executed literally).

5.1 The Randomness Expansion Protocol

Our randomness expansion protocol, Protocol 1, is described in Figure 2. The protocol is parametrized by a security parameter λ and a number of rounds N . The other parameters, the error tolerance parameter $\gamma \geq 0$ and the testing parameter $q \in (0, 1]$, are assumed to be specified as a function of λ and N . For intuition, γ can be thought of as a small constant and q as a parameter that scales as $\text{poly}(\lambda)/N$.

At the start of the protocol, the verifier executes $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$ to obtain the public key k and trapdoor t_k for a pair of functions $\{f_{k,b} : \mathcal{X} \rightarrow \mathcal{D}_{\mathcal{Y}}\}_{b \in \{0,1\}}$ from the NTCF family (see Definition 3.1). The verifier sends the public key k to the prover and keeps the associated trapdoor private.

In each of the N rounds of the protocol, the prover is first required to provide a value $y \in \mathcal{Y}$. For each $b \in \{0, 1\}$, the verifier uses the trapdoor to compute $\hat{x}_b \leftarrow \text{Inv}_{\mathcal{F}}(t_k, b, y)$. (If the inversion procedure fails, then the verifier requests another sample from the prover.) For convenience, introduce a set

$$\hat{G}_y = G_{k,0,x_0} \cap G_{k,1,x_1}, \quad (45)$$

where for $b \in \{0, 1\}$ the set G_{k,b,x_b} is defined in 4.(a) of Definition 3.1. The verifier then chooses a round type $G \in \{0, 1\}$ according to a biased distribution: either a *test round*, $G = 0$, chosen with probability $\Pr(G = 0) = q$, or a *generation round*, $G = 1$, chosen with the remaining probability $\Pr(G = 1) = 1 - q$. The former type of round is less frequent, as the parameter q will eventually be set to a very small value that goes to 0 as the number of rounds of the protocol goes to infinity. The prover is not told the round type.

Depending on the round type, the verifier chooses a challenge $C \in \{0, 1\}$ that she sends to the prover. In the case of a test round, the challenge is chosen uniformly at random; in the case of a generation round, the challenge is always $C = 1$. In case $C = 0$ the prover is asked to return a pair $(u, d) \in \{0, 1\} \times \{0, 1\}^w$. The pair is called valid if $u = d \cdot (J(\hat{x}_0) \oplus J(\hat{x}_1))$ and $d \in \hat{G}_y$, where the function J is as in 4.(b) of Definition 3.1. If $d \in \hat{G}_y$, then the verifier sets a decision bit $W = 1$ if the answer is valid, and $W = 0$ if not. If $d \notin \hat{G}_y$, then the verifier sets the decision bit $W \in \{0, 1\}$ uniformly at random.⁵ In case $C = 1$, the prover should return a pair $(b, x) \in \{0, 1\} \times \mathcal{X}$. The pair is called valid if $\text{CHK}_{\mathcal{F}}(k, b, x, y) = 1$. The verifier sets a decision bit $W = 1$ in case the pair is valid, and $W = 0$ otherwise. The set of valid pairs on challenge $C = c \in \{0, 1\}$ is denoted $V_{y,c}$.

After each test round the verifier samples a fresh $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$ and communicates the new public key k to the prover.

At the end of the protocol, the verifier computes the fraction of test rounds in which the decision bit has been set to 1. If this fraction is smaller than $(1 - \gamma)$, then the verifier aborts. Otherwise, the verifier returns the concatenation of the bits b obtained from the prover in generation rounds. (These bits are recorded in the verifier's output string $O_1 \cdots O_N$, such that $O_i = 0$ whenever the round is a test round.)

5.2 The Simplified Protocol

For purposes of analysis only, we introduce a simplified variant of Protocol 1, which is specified in Figure 3. We call it the *simplified protocol*, or Protocol 2. The protocol is very similar to the randomness expansion protocol described in Figure 2, except that the prover's answers and the verifier's checks are simplified, and in test rounds there is an additional challenge bit $T \in \{0, 1\}$. This new challenge asks the prover to perform a projective measurement on its private space that indicates whether the state lies in a "good subspace" (indicated by an outcome $K = 0$) or in the complementary "bad subspace" (outcome $K = 1$). The "good" and "bad" subspaces represent portions of space where the device's other two measurements, M and Π , are anti-aligned and aligned, respectively; see the definition of a simplified device in Section 6.2 for details.

For the case of a challenge $C = 0$, in Protocol 1 the prover returns an equation (u, d) . In the simplified protocol, the prover returns a single bit $e \in \{0, 1\}$ that is meant to directly indicate the verifier's decision (i.e., the bit W). If, moreover $T = 1$, then the prover is required to reply with an additional bit $k \in \{0, 1\}$. In this case, the verifier makes the decision to accept, i.e., sets $W = 1$, if and only if $e = 1$ and $k = 0$. For the case of a challenge $C = 1$, in Protocol 1 the prover returns a

⁵This choice is made for technical reasons that have to do with the definition of the adaptive hardcore bit property; see Section 7 and the proof of Proposition 7.4 for details.

Let λ be a security parameter, $N \geq 1$ a number of rounds, and $\gamma, q > 0$ functions of λ and N . Let $\mathcal{F}$ be an NTCF family.

At the start of the protocol, the verifier communicates N to the prover. In addition, the verifier samples an initial key $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$, sends k to the prover and keeps the trapdoor information t_k private.

- (1) For $i = 1, \dots, N$:
 - (a) The prover returns a $y \in \mathcal{Y}$ to the verifier. For $b \in \{0, 1\}$ the verifier uses the trapdoor to compute $\hat{x}_b \leftarrow \text{INV}_{\mathcal{F}}(t_k, b, y)$.
 - (b) The verifier selects a round type $G_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter q : $\Pr(G_i = 0) = q$ and $\Pr(G_i = 1) = 1 - q$. In case $G_i = 0$ (*test round*), she chooses a challenge $C_i \in \{0, 1\}$ uniformly at random. In case $G_i = 1$ (*generation round*), she sets $C_i = 1$. The verifier keeps G_i private and sends C_i to the prover.
 - (i) In case $C_i = 0$, the prover returns $(u, d) \in \{0, 1\} \times \{0, 1\}^w$. If $d \notin \hat{G}_y$, the set defined in (45), the verifier sets W to a uniformly random bit. Otherwise, the verifier sets $W = 1$ if $d \cdot (J(\hat{x}_0) \oplus J(\hat{x}_1)) = u$ and $W = 0$ if not.
 - (ii) In case $C_i = 1$ the prover returns $(b, x) \in \{0, 1\} \times \mathcal{X}$. The verifier sets W as the value returned by $\text{CHK}_{\mathcal{F}}(k, b, x, y)$.
 - (c) In case $G_i = 1$, the verifier sets $O_i = b$. In case $G_i = 0$, she sets $W_i = W$.
 - (d) In case $G_i = 0$, the verifier samples a new key $(k, t_k) \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$. She sends k to the prover and keeps the trapdoor information t_k private. This key will be used until the next test round is completed.
- (2) If $\sum_{i: G_i=0} W_i < (1 - \gamma)qN$, then the verifier aborts. Otherwise, she returns the string O obtained by concatenating the bits O_i for all $i \in \{1, \dots, N\}$ such that $G_i = 1$.

Fig. 2. The randomness expansion protocol, Protocol 1. See Definition 3.1 for notation associated with the NTCF family $\mathcal{F}$.

pair (b, x) . In the simplified protocol the prover returns a value $v \in \{0, 1, 2\}$ that is such that $v = b$ in case (b, x) is valid, i.e., $(b, x) \in V_{y,1}$, and $v = 2$ otherwise.

Note that this “honest” behavior for the prover is not necessarily efficient. Moreover, it is easy for a “malicious” prover to succeed in Protocol 2, e.g., by always returning $u = 1$ (valid equation), $k = 0$ (good subspace) and $v \in \{0, 1\}$ (valid pre-image). Our analysis will not consider arbitrary provers in Protocol 2, but instead provers whose measurements satisfy certain constraints that arise from the analysis of Protocol 1. For such provers, it will be impossible to succeed in the simplified protocol without generating randomness. Further details are given in Section 7.

5.3 Completeness

We describe the intended behavior for the prover in Protocol 1. Fix an NTCF family $\mathcal{F}$ and a key $k \in \mathcal{K}_{\mathcal{F}}$. In each round, the “honest” prover performs the following actions:

- (1) The prover executes the efficient procedure $\text{SAMP}_{\mathcal{F}}$ in superposition to obtain the state

$$|\psi^{(1)}\rangle = \frac{1}{\sqrt{|\mathcal{X}|}} \sum_{x \in \mathcal{X}, y \in \mathcal{Y}, b \in \{0, 1\}} \sqrt{(f'_{k,b}(x))(y)} |b, x\rangle |y\rangle.$$

Let λ be a security parameter, $N \geq 1$ a number of rounds, and $\gamma, \eta, \kappa, q > 0$ functions of λ and N .

- (1) For $i = 1, \dots, N$:
 - (a) The verifier selects a round type $G_i \in \{0, 1\}$ according to a Bernoulli distribution with parameter q : $\Pr(G_i = 0) = q$ and $\Pr(G_i = 1) = 1 - q$. In case $G_i = 0$ (*test round*), she chooses $C_i \in \{0, 1\}$ uniformly at random and $T_i \in \{0, 1\}$ such that $\Pr(T_i = 0) = 1 - \kappa$ and $\Pr(T_i = 1) = \kappa$. In case $G_i = 1$ (*generation round*), she sets $C_i = 1$ and $T_i = 0$. The verifier keeps G_i private and sends (C_i, T_i) to the prover.
 - (i) In case $C_i = 0$, the prover returns $e \in \{0, 1\}$. If $T_i = 1$, then the prover in addition reports $k \in \{0, 1\}$.⁶ If $T_i = 0$, then the verifier sets $W_i = e$. If $T_i = 1$, then the verifier sets $W_i = e(1 - \kappa)$.
 - (ii) In case $C_i = 1$, the prover returns $v \in \{0, 1, 2\}$. The verifier sets $O_i = v$ and $W_i = 1_{v \in \{0, 1\}}$.
- (2) If $\sum_{i: G_i=0 \wedge T_i=1} W_i < (1 - \frac{\gamma}{\kappa} - \eta)\kappa q N$, then the verifier rejects the interaction. Otherwise, she returns the string O obtained by concatenating the bits O_i for all $i \in \{1, \dots, N\}$ such that $G_i = 1$.

Fig. 3. The simplified protocol, Protocol 2.

- (2) The prover measures the last register to obtain an $y \in \mathcal{Y}$. Using item 2, from the definition of an NTCF, the prover's re-normalized post-measurement state is

$$|\psi^{(2)}\rangle = \frac{1}{\sqrt{2}}(|0, x_0\rangle + |1, x_1\rangle)|y\rangle,$$

where for $b \in \{0, 1\}$, $x_b = \text{INV}_{\mathcal{F}}(t_k, b, y)$.

- (a) In case $C_i = 0$, the prover evaluates the function J on the second register, containing x_b , and then applies a Hadamard transform to all $w+1$ qubits in the first two registers. Tracing out the register that contains y , this yields the state

$$\begin{aligned} |\psi^{(3)}\rangle &= 2^{-\frac{w+2}{2}} \sum_{d, b, u} (-1)^{d \cdot J(x_b) \oplus ub} |u\rangle |d\rangle \\ &= 2^{-\frac{w}{2}} \sum_{d \in \{0, 1\}^w} (-1)^{d \cdot J(x_0)} |d \cdot (J(x_0) \oplus J(x_1))\rangle |d\rangle. \end{aligned}$$

The prover measures both registers to obtain an (u, d) that it sends back to the verifier.

- (b) In case $C_i = 1$, the prover measures the first two registers of $|\psi^{(2)}\rangle$ in the computational basis and returns the outcome (b, x_b) to the verifier.

LEMMA 5.1. *For any λ and $k \leftarrow \text{GEN}_{\mathcal{F}}(1^\lambda)$, the strategy for the honest prover (on input k) in one round of the protocol can be implemented in time polynomial in λ and is accepted with probability negligibly close to 1.*

PROOF. Both efficiency and correctness of the prover follow from the definition of an NTCF (Definition 3.1). The prover fails only if it obtains an outcome $d \notin \hat{G}_y$, which by item 4(a) in the definition happens with negligible probability. $\square$

⁶The bit k should not be confused with the public key k for the NTCF that is used in Protocol 1. In Protocol 2, there is no NTCF and no key.

6 DEVICES

We model an arbitrary prover in the randomness expansion protocol (Protocol 1 in Figure 2) as a *device* that implements the actions of the prover: The device first returns an $y \in \mathcal{Y}$; then, depending on the challenge $C \in \{0, 1\}$, it either returns an equation (u, d) (case $C = 0$) or a candidate pre-image (b, x) (case $C = 1$). For simplicity, we assume that the device makes the same set of measurements in each round of the protocol. This is without loss of generality, as we allow the state of the device to change from one round to the next; in particular, the device is allowed to use a quantum memory as a control register for the measurements.

In Section 6.1, we introduce our notation for modeling provers in Protocol 1 as devices. In Section 6.2, we consider a simplified form of device that is appropriate for modeling a prover in the simplified protocol, Protocol 2. In Section 7, we give a reduction showing how to associate a specific simplified device to any computationally efficient general device, such that the randomness generation properties of the two devices can be related to each other (this is done in Section 8).

For the remainder of this section, we fix an NTCF family $\mathcal{F}$ satisfying the conditions of Definition 3.1 and use notation introduced in the definition.

6.1 General Devices

The following notion of device models the behavior of an arbitrary prover in the randomness expansion protocol, Protocol 1 (Figure 2).

Definition 6.1. Given $k \in \mathcal{K}_{\mathcal{F}}$, a device $D = (\phi, \Pi, M)$ (implicitly, compatible with k) is specified by the following:

- (1) A normalized density $\phi \in \text{Pos}(\mathcal{H}_D \otimes \mathcal{H}_Y)$. Here, $\mathcal{H}_D$ is an arbitrary space private to the device, and $\mathcal{H}_Y$ is a space of the same dimension as the cardinality of the set $\mathcal{Y}$, also private to the device. For every $y \in \mathcal{Y}$, define

$$\phi_y = (\text{Id}_D \otimes \langle y|_Y) \phi (\text{Id}_D \otimes |y\rangle_Y) \in \text{Pos}(\mathcal{H}_D).$$

Note that ϕ_y is sub-normalized, and $\sum_{y \in \mathcal{Y}} \text{Tr}(\phi_y) = \text{Tr}(\phi) = 1$.

- (2) For every $y \in \mathcal{Y}$, a projective measurement $\{M_y^{(u,d)}\}$ on $\mathcal{H}_D$, with outcomes $(u, d) \in \{0, 1\} \times \{0, 1\}^w$.
- (3) For every $y \in \mathcal{Y}$, a projective measurement $\{\Pi_y^{(b,x)}\}$ on $\mathcal{H}_D$, with outcomes $(b, x) \in \{0, 1\} \times \mathcal{X}$. For each y , this measurement has two designated outcomes $(0, x_0)$ and $(1, x_1)$, which are the answers that are accepted on challenge $C = 1$ in the protocol; recall that we use the notation $V_{y,1}$ for this set. For $b \in \{0, 1\}$, we use the shorthand $\Pi_y^b = \Pi_y^{(b,x_b)}$, $\Pi_y = \Pi_y^0 + \Pi_y^1$, and $\Pi_y^2 = \text{Id} - \Pi_y^0 - \Pi_y^1$.

By Naimark's theorem, up to increasing the dimension of $\mathcal{H}_D$ the assumption that $\{\Pi_y^{(b,x)}\}$ and $\{M_y^{(u,d)}\}$ are projective is without loss of generality.

We explain the connection between the notion of device in Definition 6.1 and a prover in Protocol 1. Given a device $D = (\phi, \Pi, M)$, we can define actions for the prover in Protocol 1 as follows: The prover is initialized in state ϕ . When a round of the protocol is initiated, the prover measures register Y in the computational basis and returns the outcome $y \in \mathcal{Y}$. We always assume that the prover directly measures the register, as any pre-processing unitary can be incorporated in the definition of the state ϕ . When sent challenge $C = 0$ (respectively, $C = 1$), the prover measures register D using the device's projective measurement $\{M_y^{(u,d)}\}$ (respectively, $\{\Pi_y^{(b,x)}\}$), and returns the outcome to the verifier.

Definition 6.2. We say that a device $D = (\phi, \Pi, M)$ is *efficient* if

- (1) There is a uniformly generated family of polynomial-size circuits that prepare the state ϕ given the NTCF key k as input;
- (2) For every $y \in \mathcal{Y}$, the measurements $\{M_y^{(u,d)}\}$ and $\{\Pi_y^{(b,x)}\}$ can be implemented by polynomial-size circuits.

Using the definition of an NTCF family (Definition 3.1), it is straightforward to verify that the device associated with the “honest” prover described in Section 5.3 is efficient.

We introduce notation related to the post-measurement states generated by a device in Protocol 1. An execution of Protocol 1 involves a choice of round types $g \in \{0, 1\}^N$ and challenges $c \in \{0, 1\}^N$ by the verifier and a sequence of outputs $o \in \{0, 1, 2\}^N$ computed by the verifier as a function of the answers provided by the device. Here, in case $g = 0$ (test round), we use $o \in \{0, 1\}$ to denote the outcome of the test (called W in the protocol description), and in case $g = 1$ (generation round), we use $o \in \{0, 1, 2\}$ such that $o = 2$ in case $W = 0$, and $o = O$ as recorded by the verifier in case $W = 1$. We call the tuple (g, c, o) the transcript of the protocol; it contains all the information relevant to the verifier’s final acceptance decision and to the extraction of randomness. Additional information such as the choice of NTCF key and the prover’s complete answers (including the value y) is discarded for ease of presentation. We let Acc denote the set of transcripts (g, c, o) that are accepted by the verifier in the last step of the protocol, i.e., such that $\sum_{i: g_i=0} o_i \geq (1 - \gamma)qN$.

Definition 6.3. Let $D = (\phi, \Pi, M)$ be a device. For any transcript (g, c, o) for an execution of Protocol 1 with D , let ϕ_D^{co} be the post-measurement state of the device, conditioned on having received challenges c and returned outcomes o . The joint state of the transcript and the device at the end of the N rounds (but before the verifier’s decision to abort) is

$$\phi_{\text{COD}}^{(N)} = \sum_{g, c, o} q(g, c) |c\rangle\langle c| \otimes |o\rangle\langle o| \otimes \phi_D^{co}, \quad (46)$$

where $q(g, c)$ is the probability that the sequence of round types and challenges (g, c) is chosen by the verifier in the protocol.

We write $|\phi\rangle_{\text{DE}}$ for a purification of the initial state ϕ_D of the device, with E the purifying register, and ρ_E^{co} for the post-measurement state on register E conditioned on the transcript being (c, o) .

6.2 Simplified Devices

Next, we introduce a simplified notion of device that can be used to model the actions of a prover in the simplified protocol, Protocol 2 (Figure 3).

Definition 6.4. A *simplified device* is a tuple (ϕ, Π, M, K) such that:

- (1) $\phi = \{\phi_y\}_{y \in \mathcal{Y}} \subseteq \text{Pos}(\mathcal{H}_D)$ is a family of positive semidefinite operators on an arbitrary space $\mathcal{H}_D$ such that $\sum_y \text{Tr}(\phi_y) \leq 1$;
- (2) For each $y \in \mathcal{Y}$, $\{M_y^0, M_y^1 = \text{Id} - M_y^0\}$, $\{\Pi_y^0, \Pi_y^1, \Pi_y^2 = \text{Id} - \Pi_y^0 - \Pi_y^1\}$, and $\{K_y^0, K_y^1 = \text{Id} - K_y^0\}$ are projective measurements on $\mathcal{H}_D$;
- (3) For each $y \in \mathcal{Y}$, the measurement operators K_y commute with the M_y and with the Π_y . (M_y and Π_y do not necessarily commute with each other.)

We introduce a quantity called *overlap* that measures how “incompatible” a simplified device’s measurements are. This measure is analogous to the measure of overlap used to quantify incompatibility in the derivation of entropic uncertainty relations (see, e.g., Reference [26]).

Definition 6.5. Given a simplified device $D = (\phi, \Pi, M, K)$, the *overlap* of D is

$$\Delta(D) = \max_{y \in \mathcal{Y}} \left\| K_y^0 \left(\Pi_y^0 M_y^1 \Pi_y^0 + \Pi_y^1 M_y^1 \Pi_y^1 \right) \right\|.$$

Note that the overlap only quantifies the measurement incompatibility in the “good subspace” K_y^0 .

To any simplified device $D = (\phi, \Pi, M, K)$, we associate the post-measurement states

$$\begin{aligned} \forall e \in \{0, 1\}, \quad \phi_{00}^e &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y| \otimes M_y^e \phi_y M_y^e, \\ \forall e, k \in \{0, 1\}, \quad \phi_{01}^{ek} &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y| \otimes K_y^k M_y^e \phi_y M_y^e K_y^k, \\ \forall v \in \{0, 1, 2\}, \quad \phi_1^v &= \sum_{y \in \mathcal{Y}} |y\rangle\langle y| \otimes \Pi_y^v \phi_y \Pi_y^v. \end{aligned} \quad (47)$$

A simplified device can be used in the simplified protocol in a straightforward way: Upon receipt of a challenge $C = 0$ (respectively, $C = 1$), the device first samples an $y \in \mathcal{Y}$ according to the distribution with weights $\text{Tr}(\phi_y)$. It then performs the projective measurement $\{M_y^0, M_y^1\}$ followed by, if $T = 1$, $\{K_y^0, K_y^1\}$ (respectively, $\{\Pi_y^0, \Pi_y^1, \Pi_y^2\}$) on ϕ_y , and returns the outcomes $e, k \in \{0, 1\}$ (respectively, $v \in \{0, 1, 2\}$) to the verifier.

Definition 6.6. Let $D = (\phi, \Pi, M, K)$ be a simplified device. For any transcript (g, c, t, o, k) for an execution of Protocol 2 with D , let ϕ_D^{ctok} be the post-measurement state of the device, conditioned on having received challenges (c, t) and returned outcomes (o, k) . The joint state of the transcript and the device at the end of the N rounds (but before the verifier’s decision to abort) of the protocol is

$$\phi_{\text{CTOKD}}^{(N)} = \sum_{g, c, t, o, k} q(g, c, t) |c, t\rangle\langle c, t|_{\text{CT}} \otimes |o, k\rangle\langle o, k|_{\text{OK}} \otimes \phi_D^{ctok}, \quad (48)$$

where $q(g, c, t) = q(g, c)\kappa(t)$ with $\kappa(t) = \prod_i \kappa^{t_i}(1 - \kappa)^{1-t_i}$ is the probability that the sequence of round types and challenges (g, c, t) is chosen by the verifier in the protocol.

7 SINGLE-ROUND ANALYSIS

In this section, we consider the behavior of an arbitrary device D in a single round of the randomness expansion protocol, Protocol 1 in Figure 2. Our goal is to introduce a simplified device D' such that analyzing the randomness generation properties of D' is easier than it is for D , and such that bounds on the amount of randomness generated by D' in the simplified protocol, Protocol 2 in Figure 3, imply bounds on the amount of randomness generated by D in Protocol 1. Throughout the section, we fix an NTCF family $\mathcal{F}$ (Definition 3.1) and a key $k \in \mathcal{K}_{\mathcal{F}}$ sampled according to $\text{GEN}(1^\lambda)$ for a parameter λ that plays the role of security parameter.

7.1 A Constraint on the Measurements of Any Efficient Device

We start with a lemma showing that for any efficient device $D = (\phi, \Pi, M)$, the measurements Π and M must be strongly incompatible, in the sense that if the device first measures Π , and then measures M , then it is unable to determine if the pair (u, d) returned by M corresponds to a valid pair, i.e., $(u, d) \in V_{y,0}$. Indeed, if this were the case, then the device could be used to violate the hardcore bit property (14). Recall the definition of the set $\hat{G}_y \subseteq \{0, 1\}^w$ in (45).

LEMMA 7.1. *Let $D = (\phi, \Pi, M)$ be an efficient device. Define a sub-normalized density*

$$\tilde{\phi}_{YBXD} = \sum_{y \in \mathcal{Y}} |y\rangle\langle y|_Y \otimes \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \Pi_y^{(b, x_b)} \phi_y \Pi_y^{(b, x_b)}. \quad (49)$$

Let

$$\begin{aligned}\sigma_0 &= \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \sum_{(u,d) \in V_{y,0}} |u, d\rangle\langle u, d|_U \otimes (\text{Id}_Y \otimes M_y^{(u,d)}) \tilde{\phi}_{YD}^{(b)} (\text{Id}_Y \otimes M_y^{(u,d)}), \\ \sigma_1 &= \sum_{b \in \{0,1\}} |b, x_b\rangle\langle b, x_b|_{BX} \otimes \sum_{(u,d) \notin V_{y,0}} 1_{d \in \hat{G}_y} |u, d\rangle\langle u, d|_U \otimes (\text{Id}_Y \otimes M_y^{(u,d)}) \tilde{\phi}_{YD}^{(b)} (\text{Id}_Y \otimes M_y^{(u,d)}),\end{aligned}\quad (50)$$

where $1_{d \in \hat{G}_y}$ denotes the indicator function for the event that $d \in \hat{G}_y$. Then σ_0 and σ_1 are computationally indistinguishable.

Informally, σ_0 and σ_1 in (50) are the result of performing the pre-image measurement $\{\Pi^{(b, x_b)}\}$ on ϕ_y , directly followed by an equation measurement: σ_0 is the post-measurement state associated with correct equations, and σ_1 with wrong equations. Indistinguishability of the two states follows from the hardcore bit property (14), which specifies that it is computationally infeasible to obtain a valid pre-image together with a correct equation.

PROOF. Suppose for contradiction that there exists an efficient observable O such that

$$\text{Tr}(O(\sigma_0 - \sigma_1)) \geq \mu \quad (51)$$

for some non-negligible function $\mu(\lambda)$. Consider the following efficient procedure: The procedure first prepares the state $\tilde{\phi}_{YBXD}$ in (49). This can be done efficiently by first preparing ϕ_{YD} , then measuring a $y \in \mathcal{Y}$, then applying the measurement $\{\Pi_y^{(b, x)}\}$ to ϕ_y , and returning a special abort symbol if the outcome is invalid, i.e., $\text{CHK}_{\mathcal{F}}(k, b, x, y) = 0$.

The procedure then applies the measurement $\{M_y^{(u, d)}\}$ to $\tilde{\phi}_{YBXD}$, obtaining an outcome (u, d) . At this point, conditioned on the event that $d \in \hat{G}_y$, depending on whether $(u, d) \in V_{y,0}$ or $(u, d) \notin V_{y,0}$, the procedure has either prepared σ_0 or σ_1 . Finally, the procedure measures O to obtain a bit v , and returns $(b, x, d, v \oplus u)$. This defines an efficient procedure. Moreover, using (51) it follows that the procedure violates the hardcore bit property (14). (The cases where $d \notin \hat{G}_y$ are not taken into account by the hardcore bit property, so it is sufficient to have a good distinguishing ability conditioned on $d \in \hat{G}_y$.) $\square$

7.2 Angles between Incompatible Measurements

We show a general lemma that argues about the principal angles between two binary-outcome measurements that have a certain form of incompatibility.

LEMMA 7.2. Let Π, M be two orthogonal projections on $\mathcal{H}$ and ϕ a state on $\mathcal{H}$. Let $\gamma = 1 - \text{Tr}(M\phi)$ and

$$\mu = \left| \frac{1}{2} - \text{Tr}(M\Pi\phi\Pi) - \text{Tr}(M(\text{Id} - \Pi)\phi(\text{Id} - \Pi)) \right|.$$

Let $\frac{1}{2} < \omega \leq 1$. Let K be the orthogonal projection on the direct sum of eigenspaces of $\Pi M \Pi + (\text{Id} - \Pi)M(\text{Id} - \Pi)$ with associated eigenvalue in $[1 - \omega, \omega]$. Then

$$\text{Tr}((\text{Id} - K)\phi) \leq \frac{2\mu + 10\sqrt{\gamma}}{1 - 4\omega(1 - \omega)}.$$

PROOF. Using Jordan's lemma, we find a basis of $\mathcal{H}$ in which

$$M = \oplus_j \begin{pmatrix} c_j^2 & c_j s_j \\ c_j s_j & s_j^2 \end{pmatrix} \quad \text{and} \quad \Pi = \oplus_j \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad (52)$$

where $c_j = \cos \theta_j$, $s_j = \sin \theta_j$, for some angles θ_j . There may be 1-dimensional blocks in the Jordan decomposition, but up to adding a few dimensions these can be identified with two-dimensional

blocks such that $c_j^2 \in \{0, 1\}$. Let K be the orthogonal projection on those 2-dimensional blocks such that $\min(c_j^2, s_j^2) \geq 1 - \omega$. Note that K commutes with both M and Π , but not necessarily with ϕ . It is easy to verify that this definition of K coincides with the definition given in the lemma.

Suppose first that $\gamma = 0$. Then ϕ is supported on the range of M . For any block j , let P_j be the projection on the block and $\alpha_j = \text{Tr}(P_j \phi)$. It follows from the decomposition in (52) and the definition of μ that

$$\left| \frac{1}{2} - \sum_j \alpha_j (c_j^4 + s_j^4) \right| \leq \mu. \quad (53)$$

Using that for j such that $\min(c_j^2, s_j^2) \leq 1 - \omega$, we have

$$c_j^4 + s_j^4 = 1 - 2 \max(c_j^2, s_j^2) (1 - \max(c_j^2, s_j^2)) \geq \frac{1}{2} + \left(\frac{1}{2} - 2\omega(1 - \omega) \right),$$

and $c_j^4 + s_j^4 \geq \frac{1}{2}$ always, it follows from (53) that for any $\omega > \frac{1}{2}$,

$$\text{Tr}((\text{Id} - K)\phi) \leq \frac{2\mu}{1 - 4\omega(1 - \omega)}. \quad (54)$$

Next, consider the case where $\gamma > 0$. Assume $\text{Tr}(M\phi) > 0$, as otherwise the lemma is trivial. Let $\phi' = M\phi M / \text{Tr}(M\phi)$. By the gentle measurement lemma (see, e.g., Reference [45, Lemma 9.4.1]),

$$\|\phi' - \phi\|_1 \leq 2\sqrt{\gamma}. \quad (55)$$

Using the definition of μ , it follows that

$$\left| \frac{1}{2} - \text{Tr}(M\Pi\phi'\Pi) - \text{Tr}(M(\text{Id} - \Pi)\phi'(\text{Id} - \Pi)) \right| \leq \mu + 4\sqrt{\gamma}.$$

Applying the same reasoning as for the case $\gamma = 0$ yields an analogue of (54), with ϕ' instead of ϕ on the left-hand side and $\mu + 4\sqrt{\gamma}$ instead of μ on the right-hand side. Finally, using again (55) the same bound transfers to ϕ up to an additional loss of $2\sqrt{\gamma}$. $\square$

7.3 Simulating an Efficient Device Using a Simplified Device

Recall the definitions of a simplified device (Definition 6.4) and of the overlap of a simplified device (Definition 6.5). Recall also the definition of post-measurement states $\{\phi^{co}\}$ associated with a device $D = (\phi, \Pi, M)$ given in Definition 6.3, and of post-measurement states $\{(\phi')^{ctok}\}$ associated with a simplified device $D' = (\phi', \Pi', M', K)$ given in Definition 6.6. These ensembles of states provide a means to meaningfully compare a device D and a simplified device D' . We record this in the following definition:

Definition 7.3. Let $D = (\phi, \Pi, M)$ be a device and $D' = (\phi', \Pi', M', K)$ a simplified device. We say that D' *simulates* D if for every $(c, o) \in \{0, 1\}^N \times \{0, 1, 2\}^N$ and $t = 0^N$ the states ϕ^{co} and $(\phi')^{cto}$ are identical.

The following proposition shows that any efficient device can be simulated by a simplified device whose measurements generally make an angle that is bounded away from 1. As in Lemma 7.1, the only assumption required on the efficient device is that it does not break the hardcore bit property (14).

PROPOSITION 7.4. Let $D = (\phi, \Pi, M)$ be an efficient device and $\frac{1}{2} < \omega \leq 1$. Then there is a (not necessarily efficient) simplified device $\tilde{D} = (\tilde{\phi}, \tilde{\Pi}, \tilde{M}, K)$ such that the following hold:

- (1) $\tilde{D}$ has overlap $\Delta(\tilde{D}) \leq \omega$;
- (2) The simplified device $\tilde{D}$ simulates the device D ;

(3) For any advice states $\phi' = \{\phi'_y\}$ that are independent from the key $k \in \mathcal{K}_{\mathcal{F}}$ (see Definition 2.5) it holds that

$$\sum_y \text{Tr}(K_y^1 \phi'_y) \leq C \sqrt{\sum_y \text{Tr}(\tilde{M}_y^1 \phi'_y)} + \text{negl}(\lambda), \quad (56)$$

where $C > 0$ is a constant depending only on ω .

PROOF. For each $y \in \mathcal{Y}$ let

$$\hat{M}_y = \sum_{(u,d): d \notin \hat{G}_y} M_y^{(u,d)}, \quad M_y = \sum_{(u,d) \in V_{y,0}} M_y^{(u,d)} + \frac{1}{2} \hat{M}_y,$$

and for $b \in \{0, 1\}$, $\Pi_y^b = \Pi_y^{(b, x_b)}$. By introducing an isometry $U_y : \mathcal{H}_D \rightarrow \mathcal{H}_{D'}$ into a larger space, we can embed M_y into a projection $\bar{M}_y$ such that $M_y = U_y^\dagger \bar{M}_y U_y$. For $b \in \{0, 1\}$, let $\bar{\Pi}_y^b$ be such that $U_y^\dagger \bar{\Pi}_y^b U_y = \Pi_y^b$.

The device $\tilde{D}$ is defined as follows: The device first measures an $y \in \mathcal{Y}$ exactly as D would. It then applies the isometry U_y . This defines the $\{\phi'_y\}$. Let $\phi' = \sum_y \phi'_y$, and note that $\text{Tr}(\phi') = \text{Tr}(\phi) \leq 1$.

- The measurement $\{\tilde{\Pi}_y^0, \tilde{\Pi}_y^1, \tilde{\Pi}_y^2\}$ is defined as follows: The device first coherently performs the measurement $\{\bar{\Pi}_y^{(b,x)}\}$. If an outcome $(b, x) \in V_{y,1}$ is obtained, then the device returns $v = b$. Otherwise, the device returns $v = 2$.
- The measurement $\{\tilde{M}_y^0, \tilde{M}_y^1\}$ is defined as follows: The device first performs the measurement $\{\tilde{\Pi}_y^2, \text{Id} - \tilde{\Pi}_y^2\}$. If the first outcome is obtained, then it returns a random outcome. Otherwise, it coherently performs the measurement $\{\bar{M}_y^{(u,d)}\}$. If $d \notin \hat{G}_y$, then the device returns a random outcome. Otherwise, if $(u, d) \in V_{y,0}$, then it returns a 0, and 1 if not.
- Let K_y be the projection obtained by applying Lemma 7.2 to the projections $\Pi = \bar{\Pi}_y^0$ and $M = \bar{M}_y$ and the state

$$\phi = \frac{(\bar{\Pi}_y^0 + \bar{\Pi}_y^1) \phi'_y (\bar{\Pi}_y^0 + \bar{\Pi}_y^1)}{\text{Tr}((\bar{\Pi}_y^0 + \bar{\Pi}_y^1) \phi'_y)}.$$

The measurement $\{K_y^0, K_y^1\}$ is defined by setting

$$K_y^0 = (\bar{\Pi}_y^0 + \bar{\Pi}_y^1) K_y + (\text{Id} - \bar{\Pi}_y^0 - \bar{\Pi}_y^1) \quad \text{and} \quad K_y^1 = (\bar{\Pi}_y^0 + \bar{\Pi}_y^1) (\text{Id} - K_y).$$

The first two conditions on D' claimed in the lemma follow by definition. The overlap property holds by definition of K_y^0 . For the simulation property, note that it is possible for D' to further measure the post-measurement states to locally obtain an equation, or a pre-image, as D would have; this guarantees that the post-measurement states of the two devices are identical in each round.

It remains to show the third item. It follows from computational indistinguishability of σ_0 and σ_1 shown in Lemma 7.1 that both operators have a trace that is within negligible of each other. Using the notation introduced here, and in particular the definition of $\bar{M}_y$, this implies that the difference

$$\left| \sum_{b \in \{0,1\}} \text{Tr}(\bar{M}_y \bar{\Pi}_y^b \phi'_y \bar{\Pi}_y^b) - \sum_{b \in \{0,1\}} \text{Tr}((\text{Id} - \bar{M}_y) \bar{\Pi}_y^b \phi'_y \bar{\Pi}_y^b) \right|$$

is negligible. Since the two expressions sum to $\text{Tr}((\text{Id} - \bar{\Pi}_y^2)\phi'_y)$, it follows that, letting

$$\tilde{\phi}_y = \frac{(\text{Id} - \bar{\Pi}_y^2)\phi'_y(\text{Id} - \bar{\Pi}_y^2)}{\text{Tr}((\text{Id} - \bar{\Pi}_y^2)\phi'_y)}, \quad (57)$$

we get that

$$\text{Tr}(\bar{M}_y \bar{\Pi}_y^0 \tilde{\phi}_y \bar{\Pi}_y^0) + \text{Tr}(\bar{M}_y \bar{\Pi}_y^1 \tilde{\phi}_y \bar{\Pi}_y^1)$$

is within negligible of $\frac{1}{2}$. To conclude, we apply Lemma 7.2 to the operators $\Pi = \bar{\Pi}^0$ and $M = \bar{M}_y$. The conclusion of the lemma gives that

$$\text{Tr}((\text{Id} - K_y)\tilde{\phi}_y) \leq C \sqrt{\text{Tr}((\text{Id} - \tilde{M}_y^0)\tilde{\phi}_y)} + \text{negl}(\lambda), \quad (58)$$

for some universal constant C (depending on ω). Multiplying both sides of (58) by

$$p_y = \frac{\text{Tr}((\text{Id} - \bar{\Pi}_y^2)\phi'_y)}{\sum_y \text{Tr}((\text{Id} - \bar{\Pi}_y^2)\phi'_y)},$$

summing over y and applying Jensen's inequality gives

$$\begin{aligned} & \sum_y \text{Tr}((\text{Id} - \bar{\Pi}_y^2)(\text{Id} - K_y)(\text{Id} - \bar{\Pi}_y^2)\phi'_y) \\ & \leq C \sqrt{\sum_y \text{Tr}((\text{Id} - \bar{\Pi}_y^2)(\text{Id} - \tilde{M}_y^0)(\text{Id} - \bar{\Pi}_y^2)\phi'_y)} + \text{negl}(\lambda), \end{aligned} \quad (59)$$

where we also used $\sum_y \text{Tr}((\text{Id} - \bar{\Pi}_y^2)\phi'_y) \leq 1$, since $\text{Tr}(\sum_y \phi'_y) \leq 1$. Using that by definition $(\text{Id} - \bar{\Pi}_y^2)(\text{Id} - K_y) = K_y^1$ and $(\text{Id} - \tilde{M}_y^0)$ commutes with $(\text{Id} - \bar{\Pi}_y^2)$ gives (56). $\square$

8 ACCUMULATING RANDOMNESS ACROSS MULTIPLE ROUNDS

To analyze the randomness generated by a device in the randomness expansion protocol, we proceed in two steps. First, we show that the randomness generated by the device can be related to the randomness generated by the simplified device $\tilde{D}$ that is associated to it by Proposition 7.4, when it is used as a device in the simplified protocol, Protocol 2. This is done in Section 8.1. Then, in Section 8.2, we analyze the randomness generated in a single round of the simplified protocol, and in Section 8.3, we analyze multiple rounds of the protocol.

8.1 Reduction to the Simplified Protocol

Let $D = (\phi, \Pi, M, K)$ be a simplified device. The main difference between the behavior of the simplified device and the original device it is derived from is that the simplified device (sometimes) performs an additional projective measurement $\{K^0, K^1\}$, in addition to the “equation” measurement $\{M^0, M^1\}$. (Recall that in Protocol 2, the device performs the measurement whenever the verifier sends a challenge bit $T = 1$, which happens with probability $\Pr(T = 1) = \kappa$ in the test rounds.) Informally, the measurement $\{K^0, K^1\}$ is used to detect if the state of the device is in the “good subspace” in which the device’s measurements generate randomness, as measured by the overlap $\Delta(D)$. This measurement is a conceptual tool that is not performed as part of the real protocol, but is included in the simplified protocol to facilitate the randomness generation analysis.

To lift the analysis of the randomness generated in Protocol 2 to Protocol 1, we will show that, in most test rounds of Protocol 1, the state of the device lies largely within the “good subspace”

$K = K^0$. Recall the definition of the states $\{\phi^{ctok}\}$ associated with the simplified device in Definition 6.6. Let

$$|\phi^{co}\rangle = \sum_k |\phi^{ctok}\rangle = \sum_k P_{ct}^{ok} |\phi\rangle, \quad (60)$$

where P_{ct}^{ok} is notation for the operator that corresponds to applying the device's (projective) measurement operators M , Π and K indicated by c and t , respectively, and obtaining the sequence of outcomes o and k , respectively. The fact that $|\phi^{co}\rangle$ does not depend on t is justified by the fact that $\{K, \text{Id} - K\}$ is a projective measurement.

Our goal is to bound the contribution to (60) of terms $P_{ct}^{ok} |\phi\rangle$ that correspond to a large fraction of $(\text{Id} - K)$ ("bad subspace") outcomes, i.e., such that the Hamming weight $|k|$ of the string k is large. Establishing the right bound is made delicate by the possibility of interference between the branches. We first state and prove a general lemma and then show how the lemma can be applied in our context.

LEMMA 8.1. *Let n be an integer, $0 < \kappa < 1$, and $T = (T_1, \dots, T_n)$ a sequence of independent Bernoulli random variables such that for any $t \in \{0, 1\}^n$, $\Pr(T = t) = \kappa(t) = \prod_i \kappa^{t_i} (1 - \kappa)^{1-t_i}$. Let $M = (M_1, \dots, M_n)$ and $K = (K_1, \dots, K_{|T|})$ be sequences of random variables over $\{0, 1\}$ that may be correlated between themselves and with T but satisfy that for each $i \geq 1$, M_i and T_i are independent conditioned on $(T, M, K)_{<i}$. (For an integer $i \in \{1, \dots, n\}$, we write $(T, M, K)_{<i}$ for the triple formed by the length- $(i-1)$ prefixes of T and M , and the length- $|T_{<i}|$ prefix of K .⁷)*

Assume that there is a monotone concave function $g : [0, 1] \rightarrow [0, 1]$ such that $g(0) = 0$, $g(x) \geq x$ for all $x \in [0, 1]$, and for any $i \in \{1, \dots, n\}$ and any sequences $t, m \in \{0, 1\}^{i-1}$ and $k \in \{0, 1\}^{|t|}$ it holds that

$$\Pr(K_{|t|+1} = 1 \mid (T, M, K)_{<i} = (t, m, k), T_i = 1) \leq g\left(\Pr(M_i = 0 \mid (T, M, K)_{<i} = (t, m, k), T_i = 1)\right). \quad (61)$$

Then for any $0 < \eta < 1$ there are $\kappa_0, C_0 > 0$ such that for all $0 \leq \kappa \leq \kappa_0$ and $\gamma = \kappa^{3/2}$, for all integer $n \geq 1$,

$$\sum_{t \in \{0, 1\}^n} \kappa(t) \sum_{m: |m| \geq (1-\gamma)n} \left(\sum_{k: |k| > \eta \kappa n} \sqrt{\Pr((T, M, K) = (t, m, k))} \right)^2 \leq C_0 2^{-\kappa n}. \quad (62)$$

Intuitively, the lemma holds, because the condition $|m| \geq (1 - \gamma)n$ ensures that the outcome $M_i = 0$ is fairly unlikely, in which case (61) implies that whenever $T_i = 1$ the outcome $K_j = 1$, where j is the number of nonzero entries of T in indices less or equal to i , should also be unlikely. The proof is made a little difficult by the square roots, whose presence is motivated by the application to norms of quantum states detailed later. Nevertheless, to understand the statement of the lemma, it may be useful to consider the case when all M_i (respectively, K_j) are independent and identically distributed and the square root and the square are not present. In this case, the lemma reduces to showing that if

$$\begin{aligned} \mathcal{G} &= \left\{ m \in \{0, 1\}^n : |m| \geq (1 - \gamma)n \right\}, \\ \mathcal{B}_t &= \left\{ (m, k) : m \in \{0, 1\}^n, k \in \{0, 1\}^{|t|}, |k| \geq \eta \kappa n \right\}, \end{aligned} \quad (63)$$

⁷Recall that we write $|T|$ for the Hamming weight of the string T . Here, we think of each K_j as a random variable that is correlated with the random variable M_i , where i is the index of the j th non-zero entry of T .

for $t \in \{0, 1\}^n$, then

$$\sum_{t \in \{0, 1\}^n} \kappa(t) \Pr(\mathcal{G} \wedge \mathcal{B}_t) \leq C_0 2^{-\kappa n}.$$

Note that here as in the remainder of the section, for a set $\mathcal{X}$, we often slightly abuse notation and write $\Pr(\mathcal{X})$ for $\Pr(X \in \mathcal{X})$ whenever it is clear from context which random variable is referred to.

As a first step, note that we may safely assume that $\Pr(M_i = 0) \leq \gamma + C\sqrt{\kappa}$ for some constant C , as otherwise by a Chernoff bound $\Pr(\mathcal{G}) \leq e^{-\Omega(C\kappa n)} \leq C_0 2^{-\kappa n}$ provided C is large enough. Then, using (61) it follows that $\Pr(K_j = 1) \leq g(\gamma + C\sqrt{\kappa})$, so applying the Prohorov bound [36] (see Theorem 8.3 below for an extension to martingales) with ξ there equal to K_i here, n there equal to $|t|$ here, $t = \eta$ and $v^2 = g(\gamma + C\sqrt{\kappa})$, we get that for any t ,

$$\Pr(\mathcal{B}_t) \leq e^{-\frac{\eta}{2} \operatorname{arcsinh}\left(\frac{\eta}{2g(\gamma + C\sqrt{\kappa})}\right)|t|}.$$

Using the assumption that g is monotone non-decreasing such that $g(0) = 0$, by choosing κ_0 small enough with respect to η one can ensure that the exponent

$$\frac{\eta}{2} \operatorname{arcsinh}\left(\frac{\eta}{2g(\gamma + C\sqrt{\kappa})}\right)$$

is an arbitrarily large constant C_1 . Then

$$\begin{aligned} \sum_{t \in \{0, 1\}^n} \kappa(t) \Pr(\mathcal{B}_t) &\leq \sum_{t \in \{0, 1\}^n} \kappa(t) 2^{-C_1 |t|} \\ &= \left(1 + (2^{-C_1} - 1)\kappa\right)^n \\ &\leq C_0 2^{-\kappa n}, \end{aligned}$$

where the second line uses the expression for the moment generating function for the binomial distribution and the last inequality uses $2 < e$ and holds for C_1 large enough. This completes the argument. To extend it to the general case, we use two tail bounds for martingales that replace the use of the Chernoff bound and the Prohorov bound, respectively. The first is Azuma's inequality.

THEOREM 8.2 (AZUMA'S INEQUALITY). *Let $(\xi_i, \mathcal{F}_i)_{0 \leq i \leq n}$ be a martingale difference sequence such that $\xi_0 = 0$ and $|\xi_i| \leq 1$ for each $i \in \{1, \dots, n\}$. Then for any $t \geq 0$,*

$$\Pr\left(\left|\sum_{i=1}^n \xi_i\right| \geq tn\right) \leq 2e^{-\frac{t^2}{2}n}.$$

The second is a version of the Prohorov bound for martingales.

THEOREM 8.3 (COROLLARY 2.2 IN [17]). *Let $(\xi_i, \mathcal{F}_i)_{0 \leq i \leq n}$ be a martingale difference sequence such that $\xi_0 = 0$ and $|\xi_i| \leq 1$ for each $i \in \{1, \dots, n\}$. Let*

$$X_n = \sum_{i=1}^n \xi_i \quad \text{and} \quad \langle X \rangle_n = \sum_{i=1}^n E[\xi_i^2 | \mathcal{F}_{i-1}].$$

Then for any $t \geq 0$ and $v > 0$,

$$\Pr\left(|X_n| \geq tn \text{ and } \langle X \rangle_n \leq v^2 n\right) \leq e^{-\frac{t}{2} \operatorname{arcsinh}\left(\frac{t}{2v^2}\right)n}.$$

We give the proof of Lemma 8.1.

PROOF OF LEMMA 8.1. We reduce the proof of (62) to a sequence of martingale tail bounds. Define a filtration $(\mathcal{F}_1, \dots, \mathcal{F}_i, \dots, \mathcal{F}_n)$ where $\mathcal{F}_i$ is the σ -algebra generated by $(M, T, K)_i$. Let $\mathcal{F}_{<i} = \cap_{j < i} \mathcal{F}_j$. Recall the definition of the event $\mathcal{G}$ in (63). The proof proceeds in three steps that we each formulate as a separate claim.

CLAIM 8.4 (FIRST STEP: CONDITIONAL EXPECTATIONS OF M). *Let C be a sufficiently large universal constant (C is specified in (69) in the proof). Let $\delta'_1 = \gamma + 2\sqrt{(C+1)\kappa}$. Let*

$$\mathcal{B}' = \left\{ (t, m, k) : \sum_{i=1}^n E[M_i | (T, M, K)_{<i} = (t, m, k)_{<i}] \leq (1 - \delta'_1)n \right\}. \quad (64)$$

Then it holds that

$$\sum_{m \in \mathcal{G}} \sum_t \kappa(t) \left(\sum_{k: (t, m, k) \in \mathcal{B}'} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \leq 2 \cdot 2^{-\kappa n}. \quad (65)$$

PROOF. For $i \in \{1, \dots, n\}$, let $Z_i = M_i - E[M_i | \mathcal{F}_{<i}]$ and $W_i = Z_1 + \dots + Z_i$. Then by definition the sequence $(W_1, \dots, W_n)$ is a martingale. Moreover, for any $i \geq 2$, it holds that $|W_i - W_{i-1}| = |Z_i| \leq 1$, since $M_i \in \{0, 1\}$. Applying Azuma's inequality (Theorem 8.2) it follows that for any $\delta_1 > 0$,

$$\Pr \left(\left| \sum_{i=1}^n Z_i \right| \geq \delta_1 n \right) \leq 2 e^{-\frac{\delta_1^2}{2} n}. \quad (66)$$

Let δ_1 be large enough such that the right-hand side of (66) is less than $2^{-(C+1)\kappa n}$, for some constant C to be determined below. Let $\delta'_1 = \delta_1 + \gamma$ and $\mathcal{B}'$ as in (64). Then by the Cauchy-Schwarz inequality

$$\begin{aligned} \sum_{m \in \mathcal{G}} \sum_t \kappa(t) \left(\sum_{k: (t, m, k) \in \mathcal{B}'} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \\ \leq \sum_{m \in \mathcal{G}} \sum_t \kappa(t) \left(\sum_{k: (t, m, k) \in \mathcal{B}'} \Pr((M, K) = (m, k) | T = t) \right) \left(\sum_{k: (t, m, k) \in \mathcal{B}'} 1 \right) \\ \leq \sum_{m \in \mathcal{G}} \sum_t \kappa(t) 2^{|t|} \left(\sum_{k: (t, m, k) \in \mathcal{B}'} \Pr((M, K) = (m, k) | T = t) \right), \end{aligned} \quad (67)$$

where the last inequality follows, since by definition the string k ranges over $\{0, 1\}^{|t|}$. Let $\mathcal{T}$ be the event

$$\mathcal{T} = \left\{ (t, m, k) : m \in \mathcal{G} \wedge (t, m, k) \in \mathcal{B}' \right\}.$$

Then for $(t, m, k) \in \mathcal{T}$ it holds that $|m| \geq (1 - \gamma)n$ and

$$\begin{aligned} \sum_i z_i &= \sum_i m_i - E[M_i | (T, M, K)_{<i} = (t, m, k)_{<i}] \\ &\geq (1 - \gamma)n - (1 - \delta'_1)n \\ &= (\delta'_1 - \gamma)n. \end{aligned}$$

Thus it follows from (66) and our choice of δ_1 that

$$\Pr(\mathcal{T}) = \sum_{m \in \mathcal{G}} \sum_{t, k: (t, m, k) \in \mathcal{B}'} \Pr((T, M, K) = (t, m, k)) \leq 2^{-(C+1)\kappa n}. \quad (68)$$

Finally, note that by the Chernoff bound, for C large enough,

$$\sum_{|t| \geq C\kappa n} \kappa(t) 2^{|t|} \leq 2^{-\kappa n}. \quad (69)$$

Fix C so (69) holds. Then, starting from (67), we get

$$\begin{aligned} \sum_{m \in \mathcal{G}} \sum_t \kappa(t) & \left(\sum_{k: (t, m, k) \in \mathcal{B}'} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \\ & \leq 2^{C\kappa n} \left(\sum_{m \in \mathcal{G}} \sum_{t, k: (t, m, k) \in \mathcal{B}'} \Pr((T, M, K) = (t, m, k)) \right) + 2^{-\kappa n} \\ & \leq 2 \cdot 2^{-\kappa n}, \end{aligned}$$

where the first inequality uses (69) and the second uses (68). $\square$

CLAIM 8.5 (SECOND STEP: CONDITIONAL EXPECTATIONS OF $T(1-M)$). *For any $\delta'_2 = \gamma/\kappa + \delta_2$, where δ_2 is sufficiently large compared to δ'_1 , letting*

$$\mathcal{B}'' = \left\{ (t, m, k) \notin \mathcal{B}' : \sum_i E[T_i(1 - M_i) | (T, M, K)_{<i} = (t, m, k)_{<i}] \geq \delta'_2 \kappa n \right\}, \quad (70)$$

we have that

$$\sum_{m \in \mathcal{G}} \sum_t \kappa(t) \left(\sum_{k: (t, m, k) \in \mathcal{B}''} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \leq 2^{-\kappa n}. \quad (71)$$

PROOF. For $i \in \{1, \dots, n\}$, let $Z'_i = T_i(1 - M_i) - E[T_i(1 - M_i) | \mathcal{F}_{<i}, \overline{\mathcal{B}'}]$ and $W'_i = Z'_1 + \dots + Z'_i$. Then the sequence $(W'_1, \dots, W'_n)$ is a martingale such that $|W'_i - W'_{i-1}| \leq 1$. Let

$$v_{Z'}^2 = \sum_i E[|Z'_i|^2 | \mathcal{F}_{<i}, \overline{\mathcal{B}'}].$$

For $(t, m, k) \notin \mathcal{B}'$, using that by assumption T_i is independent from M_i conditioned on $\mathcal{F}_{<i}$ and $E[T_i | \mathcal{F}_{<i}] = \kappa$ it holds that

$$\begin{aligned} v_{Z'}^2 & \leq \sum_i E[T_i(1 - M_i) | \mathcal{F}_{<i}, \overline{\mathcal{B}'}] \\ & = \kappa \sum_i E[(1 - M_i) | \mathcal{F}_{<i}, \overline{\mathcal{B}'}] \leq \delta'_1 \kappa n. \end{aligned}$$

Let $v^2 = \delta'_1 \kappa n$. Applying Theorem 8.3, for any $\delta_2 > 0$,

$$\Pr \left(\left| \sum Z'_i \right| \geq \delta_2 \kappa n \wedge v_{Z'}^2 \leq v^2 n \right) \leq e^{-\frac{1}{2} \delta_2 \kappa \operatorname{arcsinh} \left(\frac{\delta_2}{2\delta'_1} \right) n}. \quad (72)$$

Assume δ_2 chosen sufficiently large compared to δ'_1 so the right-hand side in (72) is less than $2^{-(C+1)\kappa n}$. Let $\delta'_2 = \delta_2 + \gamma/\kappa$ and $\mathcal{B}''$ as in (70). Then proceeding similarly to the end of the proof of Claim 8.4, we get (71). $\square$

CLAIM 8.6 (THIRD STEP: CONDITIONAL EXPECTATIONS OF $T(1-M)K$). *Let*

$$\mathcal{B} = \left\{ (t, m, k) : (t, m) \in \{0, 1\}^{2n}, k \in \{0, 1\}^{|t|}, |k| \geq \eta \kappa n \right\},$$

and

$$\mathcal{B}''' = \overline{\mathcal{B}'' \cup \mathcal{B}'} \cap \mathcal{B}. \quad (73)$$

Assume that $\eta = \gamma/\kappa + \delta_3$, where δ_3 is sufficiently large compared to $g(\delta'_2)$. Then

$$\sum_{m \in \mathcal{G}} \sum_t \kappa(t) \left(\sum_{k: (t, m, k) \in \mathcal{B}'''} \sqrt{\Pr((M, K) = (m, k) | T = t)} \right)^2 \leq 2^{-\kappa n}. \quad (74)$$

PROOF. We have

$$\begin{aligned}
& \sum_i \mathbb{E} \left[T_i (1 - M_i) K_{|T_{<i}|+1} \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right] \\
& \leq \sum_i \mathbb{E} \left[T_i K_{|T_{<i}|+1} \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right] \\
& = \sum_i \Pr \left(K_{|T_{<i}|+1} = 1 \middle| (T, M, K)_{<i} = (t, m, k)_{<i}, T_i = 1 \right) \Pr \left(T_i = 1 \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right) \\
& \leq \kappa \sum_i g \left(\Pr \left(M_i = 0 \middle| (T, M, K)_{<i} = (t, m, k)_{<i}, T_i = 1 \right) \right) \\
& \leq \kappa n g \left(\frac{1}{n} \sum_i \Pr \left(M_i = 0 \middle| (T, M, K)_{<i} = (t, m, k)_{<i}, T_i = 1 \right) \right). \tag{75}
\end{aligned}$$

Here, for the second line, we used $0 \leq (1 - M_i) \leq 1$; the third follows by an application of Bayes' rule; for the fourth line, we used assumption (61) and the fact that for all i ,

$$\Pr \left(T_i = 1 \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right) = \kappa; \tag{76}$$

and for the last line, we used concavity of g . For any $(t, m, k) \notin (\mathcal{B}'' \cup \mathcal{B}')$ it holds that

$$\begin{aligned}
\sum_i \mathbb{E} \left[(1 - M_i) \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right] &= \frac{1}{\kappa} \sum_i \mathbb{E} \left[T_i (1 - M_i) \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right] \\
&\leq \delta'_2 n,
\end{aligned}$$

where the equality uses (76) and the fact that T_i and M_i are independent conditioned on $(T, M, K)_{<i}$, and the inequality uses the definition of $\mathcal{B}''$. Combined with (75), we get that for any $(t, m, k) \notin (\mathcal{B}'' \cup \mathcal{B}')$,

$$\sum_i \mathbb{E} \left[T_i (1 - M_i) K_{|T_{<i}|+1} \middle| (T, M, K)_{<i} = (t, m, k)_{<i} \right] \leq g(\delta'_2) \kappa n. \tag{77}$$

For $i \in \{1, \dots, n\}$ let

$$Z_i'' = T_i (1 - M_i) K_i - \mathbb{E} \left[T_i (1 - M_i) K_i \middle| \mathcal{F}_{<i}, \overline{\mathcal{B}'' \cup \mathcal{B}'} \right]$$

and $W_i'' = Z_1'' + \dots + Z_i''$. Then the sequence $(W_1'', \dots, W_n'')$ is a martingale such that $|W_i'' - W_{i-1}''| \leq 1$ and by (77),

$$v_{Z''}^2 = \sum_i \mathbb{E} \left[|Z_i''|^2 \middle| \mathcal{F}_{<i}, \overline{\mathcal{B}'' \cup \mathcal{B}'} \right] \leq g(\delta'_2) \kappa n.$$

Applying Theorem 8.3, for any $\delta_3 > 0$, it holds that

$$\Pr \left(\left| \sum Z_i'' \right| \geq \delta_3 \kappa n \wedge \overline{\mathcal{B}'' \cup \mathcal{B}'} \right) \leq e^{-\frac{1}{2} \delta_3 \kappa \operatorname{arcsinh} \left(\frac{\delta_3}{2g(\delta'_2)} \right) n}.$$

By choosing δ_3 sufficiently large compared to $g(\delta'_2)$ the right-hand side can be made less than $2^{-(C+1)\kappa n}$. Assume further that $\delta_3 + \gamma/\kappa \leq \eta$. Let $\mathcal{B}'''$ be as in (73). Then (74) follows similarly to the proof of (65) and (71) in Claim 8.4 and Claim 8.5, respectively. $\square$

The lemma follows by combining Claim 8.4, Claim 8.5, and Claim 8.6 with the triangle inequality. $\square$

Recall the definition of the states $|\phi^{ctok}\rangle$ in Definition 6.6. For a parameter $\eta > 0$ and any $t \in \{0, 1\}^N$, let

$$|\bar{\phi}^{cto}\rangle = \sum_{k: |k| \leq \eta \kappa q N} |\phi^{ctok}\rangle, \tag{78}$$

and $\bar{\phi}^{cto}$ the sub-normalized density obtained by taking the partial trace of $|\bar{\phi}^{cto}\rangle$ over register D.

COROLLARY 8.7. *Let $D = (\phi, \Pi, M, K)$ be a simplified device such that condition (56) from Proposition 7.4 holds. Then for any $0 < \eta < 1$ there is a $\kappa_0 > 0$ such that for all $0 \leq \kappa \leq \kappa_0$ and $\gamma = \kappa^{3/2}$,*

$$\sum_{g, c \in \{0,1\}^N} q(g, c) \sum_{t \in \{0,1\}^{N-|g|}} \kappa(t) \sum_{o: (g, c, o) \in \text{Acc}} \|\phi^{co} - \bar{\phi}^{cto}\|_1 = O(2^{-\kappa q N}). \quad (79)$$

PROOF. We apply Lemma 8.1. Fix $g, c \in \{0,1\}^N$ and let $n = |\{i : c_i = 0\}|$. Let $T_1, \dots, T_n$ be independent Bernoulli random variables distributed as in the statement of Lemma 8.1. Let M and K be distributed as the measurement outcomes associated with the measurements $\{\text{Id} - M^0, \text{Id} - M^1\}$ and $\{K^0, K^1\}$ made by the device in those rounds $i \in \{0, \dots, N\}$ such that $c_i = 0$. Note that this is well-defined, since the two measurements are required to commute by Definition 6.4. Moreover, with this choice the assumption that M_i and T_i are independent conditioned on the past holds (in contrast, $K_{|T|+1}$ is correlated with T_i and with M_i).

Using (56) from Proposition 7.4, it follows that these random variables satisfy the assumptions of Lemma 8.1 for a choice of the function $g(x) = C\sqrt{x}$, for a large enough constant C . The conclusion (62) of the lemma gives (79). $\square$

We conclude with a lemma that relates the randomness in the states $\bar{\phi}^{cto}$ to randomness in the states ϕ^{ctok} , for k such that $|k| \leq \eta \kappa q N$, as these are the post-measurement states associated with the simplified device in Protocol 2. The lemma relies on the following variant of the Cauchy-Schwarz inequality:

LEMMA 8.8. *Let $\ell \geq 1$ be an integer and $|v_1\rangle, \dots, |v_\ell\rangle$ arbitrary vectors in $\mathbb{C}^d$. Then*

$$\left(\sum_{i=1}^{\ell} |v_i\rangle \right) \left(\sum_{i=1}^{\ell} |v_i\rangle \right)^\dagger \leq \ell \sum_{i=1}^{\ell} |v_i\rangle \langle v_i|.$$

PROOF. Taking the overlap with an arbitrary unit vector $|x\rangle$, the claimed inequality is equivalent to showing

$$\left| \sum_{i=1}^{\ell} \langle x | v_i \rangle \right|^2 \leq \ell \sum_{i=1}^{\ell} |\langle x | v_i \rangle|^2.$$

This follows from the Cauchy-Schwarz inequality applied to the sequences $(1, \dots, 1)$ and $(\langle x | v_1 \rangle, \dots, \langle x | v_\ell \rangle)$. $\square$

Using the lemma, we show the following:

LEMMA 8.9. *Let $D = (\phi, \Pi, M, K)$ be a simplified device, and $\bar{\phi}^{cto}$ the ensemble of states associated with D as described in (78). Then*

$$\begin{aligned} \sum_{g, c \in \{0,1\}^N} q(g, c) \sum_{\substack{t \in \{0,1\}^{N-|g|} \\ |t| \leq 2\kappa q N}} \kappa(t) \sum_{o: (g, c, o) \in \text{Acc}} \langle \bar{\phi}^{cto} \rangle_{1+\epsilon} \\ \leq 2^{O(H(\eta)\kappa q N)} \sum_{g, c \in \{0,1\}^N} q(g, c) \sum_{t \in \{0,1\}^{N-|g|}} \kappa(t) \sum_{o, k: (g, c, t, o, k) \in \text{Acc}_2} \langle \phi^{ctok} \rangle_{1+\epsilon}, \end{aligned}$$

where Acc_2 denotes the set of transcripts that are accepted by the verifier in Protocol 2.

PROOF. From the definition of $\bar{\phi}^{cto}$ in (78), applying Lemma 8.8 to the vectors $|\phi^{ctok}\rangle$ and taking the partial trace over E , we deduce that

$$\bar{\phi}^{cto} \leq \binom{\kappa q N}{\leq \eta \kappa q N} \sum_{k: |k| \leq \eta \kappa q N} \bar{\phi}^{ctok}, \quad (80)$$

where $\binom{\kappa q N}{\leq \eta \kappa q N}$ denotes the number of sequences $k \in \{0, 1\}^{|t|}$ such that $|k| \leq \eta \kappa q N$. Using standard tail bounds for the binomial distribution, this is at most $2^{O(H(\eta))\kappa q N}$. Applying the operator monotone function $\langle \cdot \rangle_{1+\varepsilon}$ on both sides of (80) and using the approximate linearity (10), we obtain

$$\langle \bar{\phi}^{cto} \rangle_{1+\varepsilon} \leq 2^{O(H(\eta))\kappa q N} \sum_{k: |k| \leq \eta \kappa q N} \langle \bar{\phi}^{ctok} \rangle_{1+\varepsilon},$$

where the factors $(1 + O(\varepsilon))$ from the approximate linearity got absorbed in the prefactor. To conclude the bound claimed in the lemma, note that the conditions that $(g, c, o) \in \text{Acc}$ and $|k| \leq \eta \kappa q N$ imply $(g, c, t, o, k) \in \text{Acc}_2$. $\square$

8.2 Randomness Accumulation in the Simplified Protocol

In this section, we consider the behavior of a simplified device $D = (\phi, \Pi, M, K)$ in a single round of Protocol 2. The following lemma shows that, provided the device has overlap $\Delta(D)$ bounded away from 1, then if the state ϕ of the device has high overlap with the projection operator M^1 , performing a measurement of $\{\Pi^0, \Pi^1, \Pi^2\}$ on ϕ necessarily perturbs the state (hence generates randomness). The proof is based on a “measurement-disturbance trade-off” from Reference [32], itself a consequence of uniform convexity for certain matrix p -norms.

LEMMA 8.10. *Let $D = (\phi, \Pi, M, K)$ be a simplified device with overlap $\Delta(D) \leq \omega$, for some $\omega < 1$. Let $0 \leq \varepsilon \leq \frac{1}{2}$ and*

$$t = \frac{\langle \phi G \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}}, \quad \text{where } G = \frac{1}{2}(\Pi^0 + \Pi^1) + \frac{1}{2}M^1K^0 \quad \text{and} \quad \phi_G = \sqrt{G}\phi\sqrt{G}. \quad (81)$$

Then

$$\frac{\langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \leq 2^{-\varepsilon \lambda_\omega(t)} + O(\varepsilon^2),$$

where the post-measurement states ϕ_1^v , $v \in \{0, 1, 2\}$, are introduced in (47), and

$$\lambda_\omega(t) = \log(e) \left(t - \frac{1}{2} - \frac{\omega}{2} \right)^2 \quad (82)$$

if $t \geq \frac{1}{2} + \frac{\omega}{2}$, and 0 otherwise.

PROOF. The proof uses ideas from Reference [32]. Let ϕ be as in the lemma and $\phi' = \sum_v \Pi^v \phi \Pi^v$. Then

$$\begin{aligned} \left\langle \sum_v \sqrt{G} \Pi^v \phi \Pi^v \sqrt{G} \right\rangle_{1+\varepsilon} &\leq \sum_v \langle \phi^{1/2} \Pi^v G \Pi^v \phi^{1/2} \rangle_{1+\varepsilon} + O(\varepsilon) \\ &\leq \left(\frac{1}{2} + \frac{\omega}{2} \right) \langle \phi^{1/2} (\Pi^0 + \Pi^1) \phi^{1/2} \rangle_{1+\varepsilon} + \frac{1}{2} \langle \phi^{1/2} \Pi^2 \phi^{1/2} \rangle_{1+\varepsilon} + O(\varepsilon) \\ &\leq \left(\frac{1}{2} + \frac{\omega}{2} \right) \langle \phi' \rangle_{1+\varepsilon} + O(\varepsilon), \end{aligned}$$

where the first and last lines use the approximate linearity relations (10), and the second line uses the definition of K and $G \leq \text{Id}$. This allows us to proceed as in the proof of Reference

[32, Theorem 5.8] to obtain

$$\langle \phi - \phi' \rangle_{1+\varepsilon} \geq 2 \left(t - \frac{1}{2} - \frac{\omega}{2} \right) \langle \phi \rangle_{1+\varepsilon} - O(\varepsilon),$$

and conclude by applying Reference [32, Proposition 4.4]. $\square$

Using Lemma 8.10, we proceed to quantify the accumulation of randomness across multiple rounds of the simplified protocol, when it is executed with a simplified device that has overlap bounded away from 1. The following proposition provides a measure of the randomness present in the transcript, conditioned on the verifier not aborting the protocol at the end, i.e., on $(g, c, t, o, k) \in \text{Acc}_2$. (To see the connection with entropy, recall the definition of the $(1 + \varepsilon)$ conditional Rényi entropy in Definition 2.9. The connection will be made precise in Section 8.3.)

PROPOSITION 8.11. *Let $D = (\phi, \Pi, M, K)$ be a simplified device such that $\Delta(D) \leq \omega$ for some $\omega < 1$. Let $0 < \varepsilon \leq \frac{1}{2}$. Let $\gamma, \eta, \kappa, q > 0$ and N an integer be parameters for an execution of Protocol 2 (Figure 3) with D . Then*

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g,c,t,o,k) \in \text{Acc}_2} q(g,c) \kappa(t) \langle \phi^{ctok} \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \right) \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta \right) - O \left(q + \frac{\varepsilon}{\kappa q} \right), \quad (83)$$

where the states ϕ^{ctok} are introduced in Definition 6.6, λ_ω is the function defined in (82), and $q(g, c)$ and $\kappa(t)$ are the distributions on N -bit strings (g, c) and t as selected by the verifier in Protocol 2.

PROOF. The proof follows a similar argument as used in Reference [32, Section 7], and we outline the main steps.

Let $t = \frac{\langle \phi_G \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}}$ be as defined in Lemma 8.10 (this t should not be confused with the string t involved in the protocol description). Recall the notation for the post-measurement states introduced in (47). After one round of Protocol 2 is executed, the post-measurement state of the device can be decomposed into three components. First, in case $G_i = 1$, which happens with probability $(1 - q)$, the round is a generation round. The randomness generated in such a round is captured by the bound from Lemma 8.10,

$$(1 - q) \left(\langle \phi_1^0 \rangle_{1+\varepsilon} + \langle \phi_1^1 \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon} \right) \leq (1 - q) \left(1 - \ln(2) \varepsilon \lambda_\omega(t) + O(\varepsilon^2) \right) \langle \phi \rangle_{1+\varepsilon}. \quad (84)$$

The second case corresponds to $G_i = 0$, which happens with probability q . In this case, for reasons that will become clear later in this proof, we weigh the “success” and “failure” components of the post-measurement state differently. For the “failure” part, we simply write

$$\frac{q}{2} \left((1 - \kappa) \langle \phi_{00}^0 \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{00} \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{01} \rangle_{1+\varepsilon} + \kappa \langle \phi_{01}^{11} \rangle_{1+\varepsilon} + \langle \phi_1^2 \rangle_{1+\varepsilon} \right). \quad (85)$$

For the “success” part, we add a weight of $2^{\frac{\varepsilon s}{\kappa q}}$, where $s = O(1)$ is a real parameter to be determined later, to the cases where $T_i = 1$:

$$\begin{aligned} & \frac{(1 - \kappa)q}{2} \left(\langle \phi_{00}^1 \rangle_{1+\varepsilon} + \langle \phi_{01}^0 \rangle_{1+\varepsilon} + \langle \phi_{01}^1 \rangle_{1+\varepsilon} \right) + \frac{\kappa q}{2} 2^{\frac{\varepsilon s}{\kappa q}} \left(\langle \phi_{01}^{10} \rangle_{1+\varepsilon} + \langle \phi_{01}^0 \rangle_{1+\varepsilon} + \langle \phi_{01}^1 \rangle_{1+\varepsilon} \right) \\ & \leq \frac{(1 - \kappa)q}{2} \left(\langle \phi_{00}^1 \rangle_{1+\varepsilon} + \langle \phi_{01}^0 \rangle_{1+\varepsilon} + \langle \phi_{01}^1 \rangle_{1+\varepsilon} \right) + \kappa q \left(1 + \ln(2) \frac{\varepsilon s}{\kappa q} + O \left(\frac{\varepsilon^2}{\kappa^2 q^2} \right) \right) t \langle \phi \rangle_{1+\varepsilon}, \end{aligned} \quad (86)$$

where the inequality follows from the definition of t . Using the first inequality in (10) and regrouping terms, the sum of the left-hand sides of (84), (85), and (86) is at most

$$(84) + (85) + (86) \leq \left(1 - \varepsilon \ln(2) \left(\lambda_\omega(t) - st + O \left(q + \frac{\varepsilon}{\kappa q} \right) \right) \right) \langle \phi \rangle_{1+\varepsilon}. \quad (87)$$

A convenient choice of s is to take the derivative $s = \lambda'_\omega(r)$ for some $r \in [0, 1]$ to be determined. With this choice, using that λ_ω is convex, it follows that $\min_{t \in [0, 1]} \lambda_\omega(t) - st = \lambda_\omega(r) - \lambda'_\omega(r)r$. By chaining the inequality (87) N times, where at each step the density ϕ is updated with the one obtained from the previous round, and using that Acc_2 contains those sequences (g, c, t, o, k) such that the number of occurrences of $(c, t, o, k) \in \{(0, 1, 1, 0), (1, *, 0, *), (1, *, 1, *)\}$ is at least $(1 - \gamma/\kappa - \eta)\kappa qN$, we obtain

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g, c, t, o, k) \in \text{Acc}_2} q(g, c) \kappa(t) \langle \phi^{ctok} \rangle_{1+\varepsilon}}{\langle \phi \rangle_{1+\varepsilon}} \right) \geq (\lambda_\omega(r) - \lambda'_\omega(r)r) + \left(1 - \frac{\gamma}{\kappa} - \eta\right) \lambda'_\omega(r) - O\left(q + \frac{\varepsilon}{\kappa q}\right),$$

with the term $(1 - \frac{\gamma}{\kappa} - \eta)\lambda'_\omega(r)$ on the right-hand side correcting for the weights $2^{\frac{\varepsilon s}{\kappa q}}$ that would appear on the left-hand side with an exponent derived from the acceptance criterion. Choosing $r = (1 - \frac{\gamma}{\kappa} - \eta)$ completes the proof. $\square$

8.3 Randomness Accumulation in the General Protocol

In this section, we combine the results obtained in the previous two sections to analyze the randomness generated in Protocol 1. The main step is given in the following proposition:

PROPOSITION 8.12. *Let $D = (\phi, \Pi, M)$ be an efficient device. Then for any $\eta > 0$ and $q > 0$ (that may be a function of N) there is a choice of parameters $0 < \kappa, \gamma < 1$ for protocol 1 such that the following hold: Let $|\phi\rangle_{DE}$ denote an arbitrary purification of ϕ_D , and $\bar{p}_{\text{COE}}$ the joint state of the verifier's choice of challenges, the outputs computed by the verifier, and the adversary's system E , restricted to transcripts that are accepted by the verifier in the protocol.⁸ Then there is a $\delta' = 2^{-\Omega(\kappa q N)}$ and a constant $C > 0$ such that for any N and δ (that may depend on N),*

$$\frac{1}{N} H_\infty^{\delta+\delta'}(O|CE)_{\bar{p}} \geq \lambda_\omega\left(1 - \kappa^{1/2} - \eta\right) - O\left(q + H(\eta)^{1/2} + \frac{1 + \log(2/\delta)}{H(\eta)^{1/2} \kappa q N}\right). \quad (88)$$

PROOF. Let $\tilde{D} = (\phi, \tilde{\Pi}, \tilde{M}, K)$ be the elementary device obtained by applying Proposition 7.4 to the device D , for a choice of $\omega = \frac{3}{4}$. Let $\tilde{\phi} = \phi^{\frac{1}{1+\varepsilon}}$, where $\varepsilon > 0$ is a small parameter to be specified later. We apply Proposition 8.11 to $\tilde{D}$, with ϕ replaced by $\tilde{\phi}$. Then (83) gives

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g, c, t, o, k) \in \text{Acc}_2} q(g, c) \kappa(t) \langle \tilde{\phi}^{ctok} \rangle_{1+\varepsilon}}{\langle \tilde{\phi} \rangle_{1+\varepsilon}} \right) \geq \lambda_\omega\left(1 - \frac{\gamma}{\kappa} - \eta\right) - O\left(q + \frac{\varepsilon}{\kappa q}\right). \quad (89)$$

Next, we apply Lemma 8.9 to obtain

$$-\frac{1}{\varepsilon N} \log \left(\frac{\sum_{(g, c, t, o): (g, c, o) \in \text{Acc}} q(g, c) \kappa(t) \langle \tilde{\phi}^{cto} \rangle_{1+\varepsilon}}{\langle \tilde{\phi} \rangle_{1+\varepsilon}} \right) \geq \lambda_\omega\left(1 - \frac{\gamma}{\kappa} - \eta\right) - O\left(H(\eta) \kappa \frac{q}{\varepsilon} + q + \frac{\varepsilon}{\kappa q}\right), \quad (90)$$

where the correction $H(\eta) \kappa \frac{q}{\varepsilon}$ comes from the exponential prefactor in the bound from Lemma 8.9. The left-hand side of the bound in Lemma 8.9 only considers those sequences such that $|t| \leq 2\kappa qN$, but adding those sequences back only incurs a negligible error $2^{-\Omega(\kappa q N)}$ (inside the logarithm), due to the Chernoff bound.

We make one ultimate re-writing step. For any fixed t , the post-measurement state $\tilde{\phi}^{cto}$ can be expressed as

$$P_N \cdots P_1 \tilde{\phi} P_1 \cdots P_N,$$

⁸The state $\bar{p}$ is sub-normalized.

where P_i is the measurement operator associated with challenge c_i and outcome o_i . Using $\langle XX^* \rangle_{1+\varepsilon} = \langle X^* X \rangle_{1+\varepsilon}$ for any X , and recalling the definition of $\tilde{\phi} = \phi^{\frac{1}{1+\varepsilon}}$,

$$\langle P_N \cdots P_1 \tilde{\phi} P_1 \cdots P_N \rangle_{1+\varepsilon} = \langle \phi^{\frac{-\varepsilon}{2(1+\varepsilon)}} \phi^{\frac{1}{2}} P_1 \cdots P_N^2 \cdots P_1 \phi^{\frac{1}{2}} \phi^{\frac{-\varepsilon}{2(1+\varepsilon)}} \rangle_{1+\varepsilon}.$$

Introduce a sub-normalized density

$$\rho_E^{cto} = \phi^{\frac{1}{2}} P_1 \cdots P_N^2 \cdots P_1 \phi^{\frac{1}{2}}$$

that corresponds to the post-measurement state of register E (recall we assumed a purification $|\phi\rangle_{DE}$ of ϕ) at the end of Protocol 1, for a given transcript (c, o) for the interaction.

We are in a position to apply Theorem 2.12, with

$$\rho_{\text{CTOE}}^o = \sum_{(g, c, t): (g, c, o) \in \text{Acc}} q(g, c) \kappa(t) |c, t\rangle\langle c, t|_{\text{CT}} \otimes |o\rangle\langle o|_O \otimes \rho_E^{cto},$$

and $\sigma_{\text{CTE}} = \sum_{(g, c, t)} q(g, c) \kappa(t) |c, t\rangle\langle c, t| \otimes \phi$. Applying the theorem and using (90) and $\langle \tilde{\phi} \rangle_{1+\varepsilon} = 1$ by definition, we get that for any $\delta > 0$,

$$\frac{1}{N} H_\infty^\delta(O|CTE)_\rho \geq \lambda_\omega \left(1 - \frac{\gamma}{\kappa} - \eta\right) - O\left(H(\eta) \kappa \frac{q}{\varepsilon} + q + \frac{\varepsilon}{\kappa q}\right) - \frac{1 + 2 \log(1/\delta)}{\varepsilon N}. \quad (91)$$

Using that the bound in (88) only considers registers C and O (the transcript) and E, by Corollary 8.7 for any choice of $0 < \eta < 1$ there is a $\kappa_0 > 0$ such that for all $0 \leq \kappa \leq \kappa_0$ and $\gamma = \kappa^{3/2}$, the bound (91) extends to a lower bound on the entropy $H_\infty^{\delta+\delta'}(O|CE)_{\bar{\rho}}$ at the cost of an additional $\delta' = O(2^{-\kappa q N})$ in the smoothing parameter.

Choose η to be an arbitrarily small constant, set $\kappa = \gamma^{2/3}$ and γ small enough so $\kappa \leq \kappa_0$. Let ε be chosen as $H(\eta)^{1/2} \kappa q$. With this choice of parameters, the term in the $O(\cdot)$ on the right-hand side of (91) is $O(q + H(\eta)^{1/2})$. $\square$

Making an appropriate choice of parameters for an execution of Protocol 1, Proposition 8.12 gives our main result.

THEOREM 8.13. *Let $\mathcal{F}$ be an NTCF family and λ a security parameter. Let N be a polynomially bounded function of λ such that $N = \Omega(\lambda^2)$. Set $q = \lambda/N$. Then there is a setting of η, γ, κ , and a $\delta = 2^{-\Omega(qN)}$ such that for any efficient prover, and side information E correlated with the prover's initial state,*

$$H_\infty^{N\delta}(O|CE)_{\bar{\rho}} \geq \xi N,$$

where $\bar{\rho}$ is the final state of the output, challenge, and adversary registers, restricted to transcripts that are accepted by the verifier in the protocol and ξ is a positive constant.⁹

Assume that an execution of $\text{GEN}(1^\lambda)$ requires $O(\lambda^r)$ bits of randomness for some constant r . (For example, for the case of our construction of an NTCF family based on LWE, we have $r = 2$.) Then an execution of the protocol using the parameters in Theorem 8.13 requires only $\text{poly}(\lambda, \log N)$ bits of randomness for the verifier to generate the key k and select the challenges. Taking N to be slightly sub-exponential in λ , e.g., $N = 2^{\sqrt{\lambda}}$, yields sub-exponential randomness expansion.

PROOF OF THEOREM 8.13. Let D be a device that is accepted with non-negligible probability in Protocol 1, where the parameters are stated in the theorem. Applying Proposition 8.12 to D with a small enough choice of η gives the result, with, e.g., $\xi = \lambda_\omega/2$. $\square$

⁹The constant ξ is at least some positive universal constant of order $1/10$.

REFERENCES

- [1] Scott Aaronson and Alex Arkhipov. 2011. The computational complexity of linear optics. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*. ACM, 333–342.
- [2] Scott Aaronson and Lijie Chen. 2017. Complexity-theoretic foundations of quantum supremacy experiments. In *Proceedings of the Computational Complexity Conference*.
- [3] Dorit Aharonov, Micahel Ben-Or, and Elad Eban. 2008. Interactive proofs for quantum computations. *Arxiv preprint arXiv:0810.5375* (2008).
- [4] Dorit Aharonov, Michael Ben-Or, Elad Eban, and Urmila Mahadev. 2017. Interactive proofs for quantum computations. *Arxiv preprint 1704.04487* (2017).
- [5] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. 2013. Learning with rounding, revisited. In *Proceedings of the Advances in Cryptology Conference (CRYPTO'13)*. Springer, 57–74.
- [6] Rotem Arnon-Friedman, Frédéric Dupuis, Omar Fawzi, Renato Renner, and Thomas Vidick. 2018. Practical device-independent quantum cryptography via entropy accumulation. *Nature Commun.* 9, 1 (2018), 459.
- [7] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando Brandao, David Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, and John Martinis. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574 (10 2019), 505–510. DOI: <https://doi.org/10.1038/s41586-019-1666-5>
- [8] Peter Bierhorst, Emanuel Knill, Scott Glancy, Yanbao Zhang, Alan Mink, Stephen Jordan, Andrea Rommal, Yi-Kai Liu, Bradley Christensen, Sae Woo Nam et al. 2018. Experimentally generated randomness certified by the impossibility of superluminal signals. *Nature* 556, 7700 (2018), 223–226.
- [9] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, John M. Martinis, and Hartmut Neven. 2016. Characterizing quantum supremacy in near-term devices. *arXiv:1608.00263* (2016).
- [10] Adam Boulund, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. 2019. On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* 15 (02 2019). DOI: <https://doi.org/10.1038/s41567-018-0318-2>
- [11] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. 2013. Classical hardness of learning with errors. In *Proceedings of the Symposium on Theory of Computing Conference (STOC'13)*, Dan Boneh, Tim Roughgarden, and Joan Feigenbaum (Eds.). ACM, 575–584. DOI: <https://doi.org/10.1145/2488608.2488680>
- [12] Anne Broadbent, Joseph F. Fitzsimons, and Elham Kashefi. 2008. Universal blind quantum computation. *Arxiv preprint arXiv:0807.4154* (2008).
- [13] Yuanmi Chen and Phong Q. Nguyen. 2011. BKZ 2.0: Better Lattice Security Estimates. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'11)*. 1–20.
- [14] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. 2018. Delegated pseudo-secret random qubit generator. *arXiv preprint arXiv:1802.08759* (2018).
- [15] Alexandru Cojocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. 2019. QFactory: classically-instructed remote secret qubits preparation. *arXiv preprint arXiv:1904.06303* (2019).
- [16] R. Colbeck. 2006. *Quantum And Relativistic Protocols For Secure Multi-Party Computation*. Ph.D. Dissertation. Trinity College, University of Cambridge. *arXiv:arXiv:0911.3814* (2006).
- [17] Xiequan Fan, Ion Grama, and Quansheng Liu. 2012. Hoeffding's inequality for supermartingales. *Stoch. Proc. Applic.* 122, 10 (2012), 3545–3559.
- [18] Joseph F. Fitzsimons and Elham Kashefi. 2017. Unconditionally verifiable blind quantum computation. *Phys. Rev. A* 96 (July 2017), 012303. DOI: <https://doi.org/10.1103/PhysRevA.96.012303>
- [19] Craig Gentry, Amit Sahai, and Brent Waters. 2013. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Proceedings of the Advances in Cryptology Conference (CRYPTO'13)*. Springer, 75–92.
- [20] Alexandru Gheorghiu and Thomas Vidick. 2019. Computationally-secure and composable remote state preparation. *arXiv preprint arXiv:1904.06320* (2019).
- [21] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. 2010. Robustness of the learning with errors assumption. In *Proceedings of the International Conference on Supercomputing*, Andrew Chi-Chih Yao (Ed.). Tsinghua University Press, 230–240.
- [22] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. 1984. A “Paradoxical” solution to the signature problem (abstract). In *Proceedings of the Advances in Cryptology Conference (Lecture Notes in Computer Science)*, G. R. Blakley and David Chaum (Eds.), Vol. 196. Springer, 467. DOI: https://doi.org/10.1007/3-540-39568-7_37
- [23] Lov Grover and Terry Rudolph. 2002. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112* (2002).
- [24] Aram W. Harrow and Ashley Montanaro. 2017. Quantum computational supremacy. *Nature* 549, 7671 (2017), 203.
- [25] Richard Lindner and Chris Peikert. 2011. Better key sizes (and attacks) for LWE-based encryption. In *Proceedings of the Cryptographers' Track at the RSA Conference*. Springer, 319–339.

- [26] Hans Maassen and Jos B. M. Uffink. 1988. Generalized entropic uncertainty relations. *Phys. Rev. Lett.* 60, 12 (1988), 1103.
- [27] Urmila Mahadev. 2017. Classical homomorphic encryption for quantum circuits. *arXiv preprint arXiv:1708.02130* (2017).
- [28] Urmila Mahadev. 2018. Classical verification of quantum computations. *arXiv preprint arXiv:1804.01082* (2018).
- [29] Daniele Micciancio and Chris Peikert. 2012. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 700–718.
- [30] Carl A. Miller and Yaoyun Shi. 2014. Universal security for randomness expansion from the spot-checking protocol. *arXiv preprint arXiv:1411.6608* (2014).
- [31] Carl A. Miller and Yaoyun Shi. 2016. Robust protocols for securely expanding randomness and distributing keys using untrusted quantum devices. *J. ACM* 63, 4 (2016), 33.
- [32] Carl A. Miller and Yaoyun Shi. 2017. Universal security for randomness expansion from the spot-checking protocol. *SIAM J. Comput.* 46, 4 (2017), 1304–1335.
- [33] Chris Peikert. 2009. Public-key cryptosystems from the worst-case shortest vector problem. In *Proceedings of the 41st ACM Symposium on Theory of Computing*. ACM, 333–342.
- [34] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. 2017. Pseudorandomness of ring-LWE for any ring and modulus. In *Proceedings of the 49th SIGACT Symposium on Theory of Computing (STOC 2017)*, Hamed Hatami, Pierre McKenzie, and Valerie King (Eds.). ACM, 461–473. DOI : <https://doi.org/10.1145/3055399.3055489>
- [35] S. Pironio, A. Acin, S. Massar, A. Boyer De La Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning et al. 2010. Random numbers certified by Bell’s theorem. *Nature* 464, 7291 (2010).
- [36] Yu V. Prokhorov. 1959. An extremal problem in probability theory. *Theor. Probab. Applic.* 4, 2 (1959), 201–203.
- [37] Oded Regev. 2009. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* 56, 6 (2009), 1–40.
- [38] B. Reichardt, F. Unger, and U. Vazirani. 2012. A classical leash for a quantum system. *Arxiv preprint arXiv:1209.0448* (2012).
- [39] Claus-Peter Schnorr. 1987. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.* 53 (1987), 201–224.
- [40] Claus-Peter Schnorr and M. Euchner. 1994. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Prog.* 66 (1994), 181–199.
- [41] Marco Tomamichel. 2015. *Quantum Information Processing with Finite Resources: Mathematical Foundations*. Vol. 5. Springer.
- [42] Marco Tomamichel, Roger Colbeck, and Renato Renner. 2009. A fully quantum asymptotic equipartition property. *IEEE Trans. Inf. Theor.* 55, 12 (2009), 5840–5847.
- [43] Umesh Vazirani and Thomas Vidick. 2011. Certifiable quantum dice: Or, true random number generation secure against quantum adversaries. In *Proceedings of the 44th Symposium on Theory of Computing (STOC’12)*. ACM, 61–76.
- [44] Umesh Vazirani and Thomas Vidick. 2014. Fully device-independent quantum key distribution. *Phys. Rev. Lett.* 113 (Sept. 2014), 140501. DOI : <https://doi.org/10.1103/PhysRevLett.113.140501>
- [45] Mark M. Wilde. 2013. *Quantum Information Theory*. Cambridge University Press. DOI : <https://doi.org/10.1017/CBO9781139525343>

Received December 2018; accepted November 2020